

# Pricing to Maximize Revenue and Welfare Simultaneously in Large Markets

Elliot Anshelevich

Koushik Kar

Shreyas Sekar

October 14, 2016

## Abstract

We study large markets with a single seller which can produce many types of goods, and many multi-minded buyers. The seller chooses posted prices for its many items, and the buyers purchase bundles to maximize their utility. For this setting, we consider the following questions: What fraction of the optimum social welfare does a revenue maximizing solution achieve? Are there pricing mechanisms which achieve both good revenue and good welfare simultaneously? To address these questions, we give efficient pricing schemes which are guaranteed to result in both good revenue and welfare, as long as the buyer valuations for the goods they desire have a nice (although reasonable) structure, e.g., that the aggregate buyer demand has a monotone hazard rate or is not too convex. We also show that our pricing schemes have implications for any pricing which achieves high revenue: specifically that even if the seller cares only about revenue, they can still ensure that their prices result in good social welfare without sacrificing profit. Our results holds for general multi-minded buyers in large markets; we also provide improved guarantees for the important special case of unit-demand buyers.

## 1 Introduction

Social Welfare and Profit <sup>1</sup> are the two canonical objectives in the extensive literature dealing with *envy-free* algorithmic pricing. The study of these two objectives, in isolation from each other, has inspired the design of novel pricing mechanisms for revenue maximization [5, 21] in a variety of interesting markets, and an equally impressive body of work on welfare maximization [13, 18, 23]. While the significance of profit and social welfare is undisputed, it is easy to overlook the fact the two objectives do not exist in a vacuum. For instance, although a monopolistic seller may only be interested in profits, myopically increasing prices while compromising on buyer welfare can lead to poor long-term revenue. This is distinctly true for large markets with repeated engagement where singularly optimizing for one objective while ignoring the other (as in the existing literature) could adversely affect the health of the marketplace [4]. Therefore, not only is it desirable to promote the design of holistic pricing solutions that optimize on both counts simultaneously, it is also crucial to gain a better understanding of how existing algorithms perform in a *bicriteria* sense. Against this backdrop, we seek to address the following questions.

*What fraction of the optimum social welfare does a revenue maximizing solution achieve? Are there pricing mechanisms which achieve both good revenue and good welfare simultaneously?*

---

<sup>1</sup>For convenience, we will use revenue and profit interchangeably in this work

Both in economics and in computer science [26], it is well understood that the goals of maximizing revenue and social welfare are often at odds with each other. Bearing this in mind, we seek to quantify the exact amount of friction between these two objectives in large markets. In particular, we are interested in understanding the surplus achieved by a profit maximizing solution, a problem that has received considerable attention in Auction theory [1, 20, 26]. The fact that we restrict our attention to the revenue end of the spectrum is motivated partly by the observation that welfare maximizing prices can result in negligible profits (see Example 3.1) even for trivial instances. However, unlike most analogous work in the theory of auctions, we are interested in understanding these trade-offs as well as designing bicriteria approximation algorithms in multi-item markets where the seller’s modus operandi involves posting prices on the individual goods. In that sense, this work is a high-level extension of the recent body of work on envy-free revenue-maximization [3, 8, 21] towards additional ambitious objectives.

## 1.1 Market Model: Item Pricing for Multi-Minded Buyers

In this work, we adopt a simple posted-pricing mechanism that captures the operation of most real-life large markets: the seller posts a single price per good, and each buyer purchases a bundle of goods that maximizes their utility. The seller controls a set  $T$  of available goods, and can produce any desired quantity  $x_t$  of a good  $t \in T$ , for which he incurs a cost of  $C_t(x_t)$ . The market consists of many, many buyers who are *multi-minded*, meaning that each buyer  $i$  has a ‘desired set of bundles of goods’: the buyer has the same value  $v_i$  for each of these bundles and under a given set of prices, purchases the bundle that maximizes her utility.

The market model that we study in this work is reasonably general. Multi-minded buyers represent a class of computationally attractive yet combinatorially non-trivial buyer valuations that have recently been featured in a number of papers [10, 27]. Perhaps, more importantly, the class strictly generalizes highly popular models such as *unit-demand* and *single-minded* valuations. Secondly, the convex production costs considered in our framework strictly generalize models with limited (or unlimited) supply, which are usually the norm in the pricing literature. As [3, 7] point out, limited supply is often too rigid for realistic, large markets where the seller may be able to increase production, albeit at a higher cost. Bicriteria algorithms notwithstanding, our work actually presents the first profit-maximization algorithms for general multi-minded buyers even with limited supply.

Our model captures several scenarios of interest wherein a typically profit maximizing seller may be driven to ensure good overall social utility. We illustrate two of them here.

1. **PEV Charging:** In a market capturing charging stations for plug-in electric vehicles, each good represents a time slot, and each buyer may desire specific (sets of) time slots based on time constraints and charging capacity. Varying demand and electricity generation costs necessitate differential pricing across time slots [6, 32].
2. **Advertising Markets:** A publisher in such a market may decide to use simultaneous posted prices to auction off a set of distinct ad-items (positions or locations on the website) to buyers interested in purchasing specific subsets of these items to reach target audiences.

## Circumventing Computational Complexity via Oblivious Guarantees

One of the challenges in essentially any non-trivial setting (including *all* the settings which we consider), is that computing profit-maximizing prices is NP-Hard. This is largely due to the

fact that the seller is not allowed to price-discriminate, i.e., it must charge the same price for each good to all the buyers, instead of having different prices for each buyer. In view of the computational barriers surrounding the optimal profit solution, a seller which cares about profit may use a variety of strategies, from approximation algorithms to heuristics. The uncertainty regarding the actual strategy adopted by the seller in turn casts aspersions on the practical significance of our goal of characterizing the social welfare at optimal-revenue solutions. One of the contributions of this work is a simple but powerful framework that allows us to completely circumvent the complexity question: *our guarantees on the social welfare do not depend on the exact details of the pricing mechanism used by the seller, and instead would hold for a wide variety of pricing mechanisms, as long as these prices achieve decent revenue guarantees.*

## Inverse Demand Functions and $\alpha$ -Strong Regularity

In order to concisely represent the large number of buyers in the market, we classify the buyers into a finite set of buyer types  $B$  such that all of the multi-minded buyers belonging to a certain type desire the same set of bundles. Then, each buyer type can be fully captured by a subset of  $2^T$  along with an *inverse demand distribution*  $\lambda_i(x)$  describing the valuations of buyers having this type. Formally, for any buyer type  $i \in B$ ,  $\lambda_i(x) = p$  implies exactly  $x$  amount of buyers of type  $i$  have a valuation of  $p$  or more for each bundle in their common desired set. Although different buyer types may have different demand functions, it is natural to assume that the valuations of all buyers are often sampled (albeit differently) from some global distribution. Because of this, we will make the assumption that the buyer valuations for every type have the same support  $[\lambda^{\min}, \lambda^{\max}]$ .


A first stab at the problem reveals that the above framework is too coarse to obtain meaningful trade-offs between social welfare and profit. Indeed, it is not hard to reason that a precise characterization of the revenue-welfare trade-offs would depend heavily on the distributions of the buyer valuations. To better understand this dependence, we study a class of inverse demand functions parameterized by a single parameter  $\alpha \in [0, 1]$  known as  $\alpha$ -strongly regular distributions.

**Definition ( $\alpha$ -Strong Regularity [16])** A buyer type  $i$  is said to have an  $\alpha$ -Strongly regular demand function ( $\alpha$ -SR) for  $\alpha \in [0, 1]$  if for any  $x_1 < x_2$ , we have  $\frac{\lambda_i(x_2)}{|\lambda'_i(x_2)|} - \frac{\lambda_i(x_1)}{|\lambda'_i(x_1)|} \leq \alpha(x_2 - x_1)$ .


$\alpha$ -Strongly regular distributions were introduced in [16] as a strict generalization of *monotone hazard rate* (MHR) distributions that smoothly interpolate between MHR ( $\alpha = 0$ ) and regular distributions ( $\alpha = 1$ ). Our main contribution is the design of mechanisms that simultaneously obtain good revenue and welfare for small  $\alpha$ , and degrade gracefully as  $\alpha$  increases. Note that even the set of  $\alpha$ -SR functions with  $\alpha = 0$ , for which we obtain the strongest results, contains a large class of important distributions, including exponential (e.g.,  $e^{-x}$ ), polynomial (e.g.,  $1 - x^2$ ), and all log-concave functions. The reader is asked to refer to Section 2 for a more detailed discussion regarding this class of distributions.

## 1.2 Our Contributions

The primary algorithmic contribution of this work is a new (profit, welfare)-bicriteria approximation for general markets with multi-minded buyers and production costs, stated below.



(a) Exact bounds for profit and welfare as a function of  $\alpha$  for unit-demand buyers. We obtain constant-factor bicriteria approximations when the demand is close to *MHR* ( $\alpha = 0$ ), and good guarantees for larger  $\alpha$  even when the demand is close to the *equal-revenue distribution* ( $\alpha = 1$ ).



(b) Actual Revenue-Welfare Curve for unit-demand buyers,  $\alpha = 0$ : The exact bicriteria guarantees lie on the trade-off curve, so that one of the two approximation factors is significantly better than the worst-case bound (point *X*). For example, when welfare is only half-optimal, the revenue factor improves from  $2e$  to 2.

**(Informal Theorem).** *We can compute in poly-time a set of item prices that guarantee a  $(\Theta(\frac{\log \Delta}{1-\alpha}))$ -approximation to the optimum profit, and a  $\Theta(\frac{1}{1-\alpha})$ -approximation for welfare, where  $\Delta$  is the ratio of the size of the largest bundle desired by any buyer to the smallest one.*

There are several exciting aspects to this result: (i) Not only do we present the *only-known bi-criteria approximation* algorithm for such general markets, but ours is also the first non-trivial profit-maximization algorithm for multi-minded buyers in the envy-free literature. (ii) Our welfare guarantees are completely independent of the bundle size ( $\Delta$ ). (iii) When buyers desire small bundles (e.g., for unit-demand valuations where all bundles are unit-size), our solution extracts a constant portion of the social welfare as revenue, as illustrated in Figure 1(a). Moreover, for the important special case of unit-demand valuations, we provide a much simpler pricing mechanism with slightly better constant approximation factors than in the theorem statement. (iv) Finally, even when buyers desire large bundles, it is reasonable to expect that in markets with similar types of goods, the various bundles are of approximately the same size, i.e.,  $\Delta$  is small. In the PEV example, one expects different electric vehicles to have similar charging capacity. This case with similar-sized bundles is a considerably non-trivial instance for which the revenue-welfare gap becomes small.

**Profit-Welfare Trade-Offs:** The approximation guarantees in the theorem statement are only the worst-case bounds derived independently for each objective. In fact, as illustrated in Figure 1(b), we prove that the two worst-case factors never occur simultaneously and the actual bicriteria bound lies on a trade-off curve, resulting in improved approximations for at least one objective, i.e., if the actual welfare is close to the worst-case guarantee, then the profit is much better than in the theorem and vice-versa.

## Social Welfare of other Revenue-Maximizing Solutions

All of the revenue guarantees in this paper are shown by comparing the profit of our solution to the optimum welfare, an approach that has strong implications towards bounding the social

welfare of other profit-maximizing solutions. Specifically, we design a simple framework using our bicriteria bounds as a black-box result and show that *any* pricing mechanism which achieves revenue better than our (efficiently computable) mechanism is guaranteed to deliver at least a  $\Theta(\frac{\log \Delta}{1-\alpha})$ -approximation to the optimum welfare. This holds whether the seller computes revenue-maximizing prices (an NP-Hard problem even for unit-demand with  $\alpha = 0$ ), or uses a more efficient mechanism. Thus, one of the main messages of this paper is that even a seller interested solely in maximizing profits can guarantee a good social welfare without sacrificing any revenue. For example, in unit-demand markets with MHR valuations, there is absolutely no excuse for such a seller to not also achieve at least a  $2e$ -approximation to the optimum social welfare irrespective of their preferred pricing mechanism.

**Technical Difficulties:** Although our large market model falls in the realm of settings where it is possible to efficiently compute social welfare maximizing prices, exploiting this for profit-maximization as in [5, 14, 21] leads to poor approximation guarantees, for e.g.,  $O(\frac{\lambda_{max}}{\lambda_{min}})$ -bounds even for unit-demand instances with no production costs. Instead, our techniques rely crucially on exploiting the structure of  $\alpha$ -strong regular functions to efficiently compute prices that compromise neither on revenue nor welfare; following this, we also develop new computational insights for characterizing profit via ascending-price procedures in settings with multi-minded buyers and production costs. Finally, in this work, we will focus solely on deterministic pricing mechanisms. While randomized mechanisms that mix between welfare and profit maximizing solutions are a theoretically fascinating apparatus, the ensuing price fluctuations and alternating buyer dissatisfaction render them unsuitable for many settings of interest [19, 28]

## Related Work: Existing Bicriteria Algorithms

The primary barrier towards designing envy-free revenue-maximizing prices — a lack of insight regarding the optimum solution — is also, surprisingly, the chief architect behind the existence of many (implicit) bi-criteria approximation algorithms in the algorithmic pricing literature. More concretely, a majority of the revenue-maximization algorithms in the literature [5, 9, 12, 14, 21, 25] achieve their approximation factors for revenue by comparing it to the optimum social welfare. Exploiting these revenue-welfare ties further, it is not hard to see that such an  $\alpha$ -approximation algorithm for revenue trivially results in a  $(\alpha, \alpha)$ -bicriteria approximation. For instance, the results from [21] and [5] immediately imply  $(\Theta(\log |B|), \Theta(\log |B|))$ -bicriteria approximation algorithms for unit-demand and unlimited supply markets respectively.

Our results improve upon the work mentioned above across multiple dimensions. In contrast to the trivial  $(\alpha, \alpha)$  type bounds in previous work, our specific focus on bicriteria approximations leads to significant improvements in social welfare without sacrificing much profit. We also remark that while specific bicriteria bounds were also provided in [3], their results only apply to the easiest version of our setting ( $\alpha = 0, \Delta = 1$ ). Indeed, we study general settings with multi-minded buyers and production costs, instead of unit-demand and single-minded valuations with limited supply, which are much more common in algorithmic pricing literature. Finally, our bounds are more nuanced due to their dependence on  $\alpha$ , and independence from  $|B|$ , e.g.,  $\Theta(\log |B|)$ -bounds are totally unacceptable for large markets.

## 1.3 Other Related Work

While (revenue, welfare)-bicriteria approximation algorithms have not specifically been tackled in the pricing literature beyond the single good case, the broader understanding of trade-offs

between the two objectives has been a prominent motif that has cropped up time and again in various forms. We first highlight a few overarching differences between our results and other work on profit-welfare trade-offs, specifically in auctions: (i) we study reasonably general combinatorial markets with multi-minded buyers and production cost functions, and not just limited-supply settings with unit-demand buyers as in other work, (ii) our results do not depend on externalities such as tuning the objective function or the addition of external buyers as in Bulow-Klemperer type theorems [2, 11, 30], and finally (iii) unlike similar (types of) results in Bayesian auctions, our pricing mechanisms are non-discriminatory, and therefore, envy-free.

Characterizing the efficiency of revenue-optimal mechanisms is an extremely fundamental question that has spurred multiple avenues of research starting with Bulow and Klemperer's seminal result [11] that adding a single buyer is usually more profitable than blindly optimizing revenue. Most pertinent to the questions posed in this work are the tight bounds on the (in)efficiency of the Myerson revenue-maximizing mechanism for single good settings appearing in [1, 26]: in particular, [26] provides welfare bounds for general single-parameter auctions as a function of the distribution of buyer valuations, as we do in this work.

Moving beyond welfare lower bounds, other researchers have adopted a more constructive approach by designing explicitly taking into account both the objectives either via bicriteria mechanisms [17, 31], or by optimizing linear combinations of revenue and welfare [28], or even characterizing the revenue-welfare Pareto curve [19]. We reiterate that all of the above papers consider simple single good settings, which are easier to characterize, and where the revenue-optimal mechanism is well understood. Moreover, in comparison to the revenue-welfare Pareto curves in [19], the implicit trade-off curves in our work are of a different nature as they are obtained for a single instance on top of the worst-case bounds. Finally, multi-objective trade-offs are quite popular in the Sponsored Search literature [4, 18, 29, 30], where the repeated engagement and the tight competition (between sellers) necessitates approximately revenue-optimal mechanisms that do not compromise on overall social welfare. Such settings can essentially be viewed as a special case of unit-demand markets.

## 2 Model and Preliminaries

Our market model comprises of a single seller controlling a set  $T$  of goods and a large number of infinitesimal buyers. The buyers can be concisely represented using a finite set of multi-minded buyer types  $B$ : for a given type  $i \in B$ , all the buyers having this type desire the same set of item bundles  $B_i \subseteq 2^T$ , and each buyer is indifferent between the bundles in  $B_i$ . Notice that when all of the desired bundles are of unit cardinality, our model reduces to the *unit-demand* case; when each buyer type desires only a single bundle ( $|B_i| = 1$ ), we get *single-minded* valuations. For general multi-minded valuations, we also assume the free disposal property. Finally, buyers belonging to the same type may hold different valuations for the same bundles, this is modeled by way of an *inverse demand function*  $\lambda_i(x)$  for every  $i \in B$ ;  $\lambda_i(x) = p$  implies that exactly  $x$  amount of buyers of type  $i$  value the bundles in  $B_i$  at valuation  $p$  or more. Given  $\lambda_i(x) = p$ , it is not hard to see that the total utility derived when  $x$  amount of buyers purchase some bundle at price  $p$  is  $u_i(x) = \int_{z=0}^x \lambda_i(z) dz$ .

The market operates according to a natural pricing mechanism with the seller posting a price  $p_t$  for each good  $t \in T$ . Buyers purchase one of the utility-maximizing bundles available to them, i.e., a buyer belonging to type  $i$  will purchase the cheapest bundle in  $B_i$  as long as its price is no larger than her valuation for the same. Therefore, if  $\bar{p}_i$  denotes the bundle in  $B_i$  with

the smallest price and  $\bar{x}_i$  is the population of buyers of this type who purchased some bundle, then  $\lambda_i(\bar{x}_i) = \bar{p}_i$ .

### Pricing Solutions, Social Welfare, and Revenue

We use  $(\vec{p}, \vec{x}, \vec{y})$  to represent the outcome of the market mechanism. Here  $\vec{p}$  is the vector of prices,  $\vec{x}$  denotes the allocation to the buyers with  $x_i$  being the total amount of good purchased by buyers of type  $i$ , and finally  $y_t$  is the total amount of good  $t \in T$  sold to the buyers. We now define the two main metrics that form the crux of this paper.

**Definition** The social welfare of a solution  $(\vec{p}, \vec{x}, \vec{y})$  is defined to be the total utility of all the buyers and the seller and therefore, is equal to the utility of the buyers minus the production cost incurred by the seller, i.e.,

$$SW(\vec{p}, \vec{x}, \vec{y}) = \sum_{i \in B} \int_{x=0}^{x_i} \lambda_i(x) dx - \sum_{t \in T} C_t(y_t).$$

Notice that the social welfare is independent of the prices, and depends only on  $(\vec{x}, \vec{y})$ .

**Definition** The (seller's) profit at the solution  $(\vec{p}, \vec{x}, \vec{y})$  is the total income due to each good in the market minus the total production cost incurred, i.e.,

$$\pi(\vec{p}, \vec{x}, \vec{y}) = \sum_{t \in T} [p_t y_t - C_t(y_t)].$$

One of the main goals of this paper is to obtain a lower bound on the social welfare of the profit maximizing solution. We will use  $(\vec{p}^{opt}, \vec{x}^{opt}, \vec{y}^{opt})$  to denote the maximum profit solution, and  $(\vec{p}^*, \vec{x}^*, \vec{y}^*)$  to denote the solution maximizing welfare. Sometimes we will also use  $SW^* = SW(\vec{p}^*, \vec{x}^*, \vec{y}^*)$  and  $\pi^{opt} = \pi(\vec{p}^{opt}, \vec{x}^{opt}, \vec{y}^{opt})$ . Thus the quantity we are interested in is  $\frac{SW(\vec{p}^*, \vec{x}^*, \vec{y}^*)}{SW(\vec{p}^{opt}, \vec{x}^{opt}, \vec{y}^{opt})}$ .

We remark that the maximum social welfare solution should ideally be represented as  $(\vec{x}^*, \vec{y}^*)$ . However, due to the following simple claim (inspired by a similar claim from [3]), we will define  $p_t^*$  as below, and let  $(\vec{p}^*, \vec{x}^*, \vec{y}^*)$  denote this specific welfare-maximizing solution. More importantly, as we show in the Appendix, such a welfare maximizing solution can be computed efficiently via a simple convex program.

**Claim 2.1.** *There exists a welfare-maximizing solution  $(\vec{p}^*, \vec{x}^*, \vec{y}^*)$  where for every good  $t \in T$ ,  $p_t^* = c_t(y_t^*)$ . (Here  $c_t$  is the derivative of  $C_t$ .)*

### Structure of the Demand and Cost Functions

In this work, we are interested in studying markets with many, many buyers and therefore, correspondingly we take both the inverse demand and production cost functions to be continuously differentiable. In addition, we also make the standard assumption that the utility function  $u_i(x)$  is concave for all  $i \in B$  and therefore its derivative  $\lambda_i(x)$  is non-increasing with  $x$ . Finally, we consider production cost functions that are *doubly convex* i.e, both  $C_t$  and its derivative  $c_t(x) = \frac{d}{dx} C_t(x)$  are convex and non-decreasing for all  $t \in T$  and further,  $C_t(0) = c_t(0) = 0$ . A number of well studied cost functions fall within our framework [7].

As mentioned previously, it is often natural to assume that the inverse demand distributions for different buyer types have the same support  $[\lambda^{min}, \lambda^{max}]$ . In fact, *all* of our results hold under the more general *uniform peak* assumption, which will be assumed for the rest of this work.

**Definition** (Uniform Peak Assumption) For every  $i \in B$ ,  $\lambda_i(0) = \lambda^{max}$ .

**$\alpha$ -SR inverse demand** (Definition 1.1) Recently,  $\alpha$ -strong regularity has gained some popularity [15, 16, 24] as an elegant characterization of the class of regular functions, which encompasses most well-studied demand distributions including polynomial ( $\lambda(x) = 1 - x^2$ ;  $\alpha = 0$ ), exponential ( $\lambda(x) = e^{-x}$ ;  $\alpha = 0$ ), power law ( $\lambda(x) = \frac{1}{\sqrt{x}}$ ;  $\alpha = \frac{1}{2}$ ), and the equal-revenue distribution ( $\lambda(x) = \frac{1}{x}$ ;  $\alpha = 1$ ). For such functions, one can interpret  $\alpha$  as a measure of the convexity of the function as larger values of  $\alpha$  imply greater convexity or alternatively as a bound on the volatility of the inverse demand  $\lambda_i(x)$  as every  $\alpha$ -SR demand function satisfies  $\frac{d}{dx} \left( \frac{\lambda(x)}{|\lambda'(x)|} \right) \leq \alpha$ . As expected, the equal-revenue distribution ( $\alpha = 1$ ) leads to the worst-case bounds for all of our results: in fact even in single good, single buyer type markets, the revenue-optimal solution for  $\alpha = 1$  only extracts a negligible fraction of the optimum social welfare. However, what is surprising (as evidenced by Figure 1(a)) is that we obtain reasonably good performance guarantees even when  $\alpha$  is larger than  $\frac{1}{2}$ . Note that even  $\alpha$ -SR functions with very small  $\alpha$  include a very large class of interesting demand functions (see above), including any log-concave function.

### 3 Warm-up: Profit and Welfare for Unit-Demand Markets

As a first step towards stating our general results in Section 5, we consider the important special case of unit-demand markets, perhaps the most popular class of valuations in the algorithmic pricing literature [9, 12, 21]. Recall that unit-demand valuations are a simple sub-class of multi-minded functions, wherein for each buyer (type)  $i \in B$ , the bundles desired by  $i$  are singleton sets. Therefore, we will overload notation, and when talking about unit-demand buyers say that  $B_i \subseteq B$  is the set of items acceptable by all of the buyers of type  $i$ . Our main result in this section is a simple pricing rule that achieves a  $(\Theta(\frac{1}{1-\alpha}), \frac{2-\alpha}{1-\alpha})$ -bicriteria approximation algorithm for revenue and welfare respectively when the buyers have  $\alpha$ -SR inverse demand functions. While we present a generalized version of this theorem in Section 5, the simple techniques from this section do not really extend to multi-minded case. In fact, we require several new gadgets and a more sophisticated pricing mechanism to achieve the generalization.

That said, our reasons for dedicating an entire section to unit-demand buyers is two-fold: (i) the algorithm presented in this section is extremely simple and the constant factors hidden by the asymptotic bound are smaller, and (ii) the unit-demand case provides a platform for us to build intuition and more importantly, discuss the various implications of our results including the profit-welfare trade-off and the ability to derive welfare bounds for the profit-maximizing solution. Before stating our main theorem, we give a simple example to highlight the poor revenue guarantees obtained by the welfare maximizing prices even for single-good markets.

**Example 3.1.** Consider a single good market with a negligible production cost function, say  $C(x) = \epsilon x$ . Obviously, there is only one buyer type, and suppose that its inverse demand is  $\lambda(x) = 1 - x$  ( $\alpha$ -SR for  $\alpha = 0$ ). It is easy to observe that at the social welfare maximizing solution, the good is priced at  $p^* = \epsilon$ , resulting in near-zero revenue. On the contrary, one can price at  $p^{opt} = \frac{1}{2}$  to obtain a constant fraction of the optimum welfare as profit.



**Theorem 3.2.** *For any unit-demand instance where buyers have  $\alpha$ -strongly regular inverse demand functions, there is a poly-time  $(\zeta, \frac{2-\alpha}{1-\alpha})$ -bicriteria approximation algorithm for revenue and social welfare respectively, where*

$$\zeta = 2\left(\frac{1}{1-\alpha}\right)^{\frac{1}{\alpha}} + \frac{\alpha}{1-\alpha} = \Theta\left(\frac{1}{1-\alpha}\right).$$

The exact guarantees for profit and welfare are illustrated in Figure 1(a).

(Algorithm) The bicriteria approximation factor is achieved by the following simple pricing mechanism

- Compute the max-welfare solution  $(\vec{p}^*, \vec{x}^*, \vec{y}^*)$ .
- For every good  $t$ , set its price  $\tilde{p}_t = \max(p_t^*, \lambda^{\max}(1-\alpha)^{\frac{1}{\alpha}})$ .

**Proof Sketch:** We prove this theorem in the Appendix, but here we give some intuition for why this produces a good approximation to both profit and welfare. Let  $\tilde{p} = \lambda^{\max}(1-\alpha)^{\frac{1}{\alpha}}$ . We begin by analyzing these types of *thresholded* pricing schemes, in which the price is simply the maximum of the optimum price  $p_t^*$  and a constant. In Lemma 3.3, we show that such solutions have nice structure; essentially we can think of buyers who purchase goods priced at  $\tilde{p}$  and those who purchase goods priced at  $p_t^*$  as separate systems.

**Lemma 3.3.** *Suppose that  $((\tilde{p})_{t \in T}, (\tilde{x})_{i \in B}, (\tilde{y})_{t \in T})$  is a pricing solution resulting from a thresholded pricing vector. Then,*

1. *The market can be clustered into two mutually disjoint sets of buyers and goods  $(B^H, T^H)$  and  $(B^L, T^L)$  so that the buyers in each cluster only purchase the goods in the same cluster and (a) for  $(i, t) \in (B^H, T^H)$ ,  $\tilde{p}_t = p_t^*$ ,  $\tilde{x}_i = x_i^*$ , and  $\tilde{y}_t = y_t^*$ ; (b) for  $(i, t) \in (B^L, T^L)$ ,  $\tilde{p}_t \geq p_t^*$ ,  $\tilde{x}_i \leq x_i^*$ , and  $c_t(\tilde{y}_t) \leq c_t(y_t^*)$ .*
2.  *$\vec{\tilde{y}}$  is a welfare-maximizing allocation vector with respect to the demand vector  $\vec{\tilde{x}}$ .*

We then prove that due to our choice of  $\tilde{p}$  the following two bounds hold:

$$SW(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}}) \leq (2\left(\frac{1}{1-\alpha}\right)^{\frac{1}{\alpha}} - 1)\pi(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}}),$$

and

$$SW(\vec{p}^*, \vec{x}^*, \vec{y}^*) - SW(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}}) \leq \frac{1}{1-\alpha}\pi(\vec{\tilde{p}}, \vec{\tilde{x}}, \vec{\tilde{y}}).$$

Once we show those lower bounds on the profit of our pricing scheme, we apply the following simple claim to finish the proof of the theorem. While the difficult parts of our argument lie in proving the above bounds, we state the claim here since, despite its technical simplicity, we believe that this claim and our general approach may find application in other settings.

**Claim 3.4.** *Suppose that for some solution  $s$  we have that  $SW(s) \leq c_1\pi(s)$ , and  $SW^* - SW(s) \leq c_2\pi(s)$ . Then,  $s$  is a  $(c_1 + c_2, c_2 + 1)$ -bicriteria approximation to  $(\text{profit}, \text{welfare})$ ; moreover  $SW^* \leq (c_1 + c_2)\pi(s)$ . ■*

Henceforth, we will unambiguously use  $\zeta$  to denote the exact approximation factor appearing in the statement of the above theorem. Interestingly, in the proof of Theorem 3.2, the revenue guarantee of  $\zeta$  is shown with respect to the optimum social welfare, which in turn has important consequences for other profit-maximizing solutions (see Section 4). For now, we just state this revenue guarantee formally.

**Corollary 3.5.** *The profit obtained by the pricing mechanism in Theorem 3.2 is always within a factor  $\zeta$  of the optimum social welfare.*

### 3.0.1 Profit-Welfare Trade-offs

We further exploit the close ties between revenue and social welfare, and present a revenue-welfare trade-off that improves upon the bicriteria bound in Theorem 3.2 by showing that at least one of revenue or welfare is better than the factor guaranteed by the theorem. The bounds in Theorem 3.2 are actually somewhat misleading as they represent the worst-case bound for each objective, which is derived independently from the other objective by simply bounding the worst-case revenue (or welfare) over all instances. However, the worst-case performance for revenue ( $\zeta$ ) and the worst-case performance for social welfare ( $\frac{2-\alpha}{1-\alpha}$ ) need not and as we show, *do not* occur for the same instance. As a matter of fact, for a given instance, if the actual welfare obtained is close to the guarantee provided in Theorem 3.2, the large gap between welfare and revenue as in Figure 1(a) completely vanishes and the approximation factors coincide.

We first state the main structural claim that enables this trade-off.

**Claim 3.6.** *Suppose that a pricing algorithm  $Alg$  satisfies  $SW(Alg) \leq c_1\pi(Alg)$ ,  $SW^* - SW(Alg) \leq c_2\pi(Alg)$ . Then for every instance there exists some  $1 \leq c \leq c_2 + 1$  such that  $Alg$  is a bicriteria approximation  $(\min(cc_1, \frac{cc_2}{c-1}), c)$  for revenue and welfare respectively.*

Unfortunately, the designer has no control over the factor  $c$  as its exact value depends on the particular instance. Applying this claim to Theorem 3.2 yields the following corollary.

**Corollary 3.7.** *For every given instance, there exists a constant  $1 \leq c \leq \frac{2-\alpha}{1-\alpha}$  so that the solution returned by the prices of Theorem 3.2 has a social welfare that is within a factor  $c$  of the optimum welfare and such that*

$$\pi^{opt} \leq SW^* \leq \min\left(\frac{c}{c-1} \cdot \frac{1}{1-\alpha}, \zeta\right)\pi(Alg).$$

For example, for MHR demand functions the statement of Theorem 3.2 makes it seem that this pricing scheme may return a solution which is a 2e-approximation for revenue and a 2-approximation for welfare. Corollary 3.7 points out that the actual results are far *better*. As Figure 1(b) illustrates for MHR demand, tradeoffs between revenue and welfare actually guarantee that when revenue is far from optimum, welfare is very high, and vice versa.

## 4 Consequences for Solutions with High Revenue

In Sections 3 and 5, we give efficient pricing mechanisms which simultaneously achieve good approximations for both revenue and welfare. Consider, however, a seller whose main priority is to simply maximize profits. This seller may choose to use a different pricing mechanism with better revenue guarantees than the ones offered in this paper. For example, the seller may choose prices which are guaranteed to come closer to achieving optimum revenue (these are efficiently computable for unit-demand settings [3, 22] under certain additional assumptions), or even use a large amount of resources to solve the intractable problem of actually computing prices  $p^{opt}$  which yield the highest possible revenue. One of the main messages of our paper is as follows:

*No matter what pricing mechanism the seller uses to optimize revenue, they can instead use a pricing mechanism which guarantees at least  $1/\zeta$  fraction of the optimum welfare, without sacrificing any revenue.*

In this section, we use a simple albeit highly general framework to derive results of this form. Although this framework is defined in terms of our model, it is actually general and applies across a wide variety of markets.

Recall that  $SW^*$  denotes the optimum social welfare and  $\pi^{opt}$  denotes the maximum achievable profit. Given a pricing algorithm  $Alg$ , we will refer to the social welfare of the solution returned by  $Alg$  using  $SW(Alg)$ , and its profit by  $\pi(Alg)$ . Consider an arbitrary profit maximization algorithm  $Alg$  that achieves a good approximation with respect to  $\pi^{opt}$ . How do we go about characterizing the social welfare at these solutions? The following theorem uses an existing profit-maximization algorithm whose guarantees hold with respect to the optimum welfare as a black-box to bound the welfare due to  $Alg$ .

**Theorem 4.1.** *Consider a benchmark profit maximization algorithm  $Alg^b$  whose profit  $\pi(Alg^b)$  is always within a  $c$  factor of  $SW^*$  for some  $c \geq 1$ . Consider any other pricing algorithm  $Alg$  for the same class of valuations that obtains at least as much profit as guaranteed by  $Alg^b$  on all instances. Then the social welfare obtained by  $Alg$  is at least a factor  $\frac{1}{c}$  times that of the optimum welfare.*

The proof simply follows from the fact that  $SW(Alg) \geq \pi(Alg) \geq \pi(Alg^b) \geq \frac{SW^*}{c}$ .

## Implications for Unit-Demand Markets

We now apply the framework provided by Theorem 4.1 for unit-demand using our result from Theorem 3.2 as a black-box.

**Claim 4.2.** *Let  $Alg$  be any algorithm for unit-demand markets that obtains at least as much profit as guaranteed by the algorithm of Theorem 3.2 on all instances. Then the social welfare obtained by  $Alg$  is at least a factor  $\frac{1}{\zeta}$  times that of the optimum welfare.*

Thus, consider the case when a seller is using any arbitrary pricing mechanism  $Alg$ , with the main goal being to maximize profit. By simply computing the revenue given by our pricing schemes from Section 3, and then choosing the one which guarantees better revenue (i.e., choosing between  $Alg$  and our pricing scheme), we form a new pricing algorithm which does not sacrifice any revenue compared to  $Alg$ , and due to the above theorem, is *also* guaranteed to have good social welfare. Moreover, for sellers who are able to compute the prices which result in absolute maximum revenue, the following is a trivial consequence of the above theorem:

**Corollary 4.3.** *The ratio of the optimum social welfare  $SW^*$  to the social welfare at the maximum profit solution  $SW(p^{opt}, x^{opt}, y^{opt})$  is at most  $\zeta = \Theta\left(\frac{1}{1-\alpha}\right)$ .*

For example, for MHR demand functions ( $\alpha = 0$ ), this implies that even for sellers who only care about profits, there is essentially no excuse not to also guarantee at least  $1/2e$  of the optimum social welfare. Thus, for the settings we consider, one can strive for truly high revenue, without sacrificing much in welfare.

## 5 Multi-Minded Buyers

We now move on to our most general case with multi-minded buyers, wherein every buyer wishes to purchase one bundle from a desired set (of bundles). We use  $\ell^{max}$  to denote the cardinality of the maximum sized bundle desired by any buyer type, and  $\ell^{min}$  for the minimum sized bundle.

The main result in this section is a bicriteria approximation algorithm that extends our results for the unit-demand case. Our algorithm still achieves a  $\Theta(\frac{1}{1-\alpha})$ -approximation to the optimum welfare; as for profit, we obtain a  $\Theta(\frac{1}{1-\alpha})$  bound further discounted by a  $\log(\Delta)$  factor, where  $\Delta = \frac{\ell^{max}}{\ell^{min}}$ .

Moreover, as in the previous section, our profit bound is obtained in terms of the optimum social welfare, which allows the consequences mentioned in Section 4 as well as an analog of Corollary 3.7 to hold showing that the true lower bounds are better than the theorem states.

**Theorem 5.1.** *For any given instance with multi-minded buyers, there exists a poly-time  $(\Theta(\frac{\log(\Delta)}{1-\alpha}), \Theta(\frac{1}{1-\alpha}))$ -bicriteria approximation algorithm for profit and welfare respectively.*

**Proof Sketch:** The proof for the general case with multi-minded buyers is quite involved (see Appendix B), so here we provide a high level overview of the various ingredients that combine to form the proof.

*Step 1: Benchmark Solution.* The first step involves defining a benchmark solution  $(\vec{x}^b, \vec{y}^b)$  whose social welfare, as in the unit-demand case, is at most a  $\frac{2-\alpha}{1-\alpha}$  factor away from  $SW^*$ . Specifically, define  $\tilde{p} = \lambda^{max}(1-\alpha)^{\frac{1}{\alpha}}$ . Now, the benchmark demand vector  $\vec{x}^b$  is defined as follows: for each buyer  $i$ ,  $x_i^b = \min\{x_i^*, \lambda^{-1}(\tilde{p})\}$ . The allocation vector  $\vec{y}^b$  is simply a scaled-down version of the optimum allocation vector, i.e., suppose that  $x_i^*(S)$  is amount of bundle  $S$  consumed by buyer type  $i$  in  $(\vec{p}^*, \vec{x}^*, \vec{y}^*)$ . Then, the amount of bundle  $S$  consumed by  $i$  in the benchmark solution  $x_i^b(S) := \frac{x_i^b}{x_i^*} x_i^*(S)$ . This consumption pattern gives rise to the allocation vector  $\vec{y}^b$  on the goods. Let  $SW^b$  denote the social welfare of this benchmark solution.

In the unit-demand case, we were able to identify a suitable price vector to actually implement this benchmark solution and extract a good fraction of its welfare as revenue. However, this is no longer possible for the general case, as an analogous pricing solution would involve personalized payments for each buyer and thus may not be realizable using simple item pricing. For convenience, we define  $\pi^b = \sum_{i \in B} \lambda_i(x_i^b) x_i^b - C(\vec{y}^b)$  to be the profit due to the personalized payments. Observe that if we allow for discriminatory prices, then the revenue maximization becomes almost trivial.

Despite its lack of implementability, the benchmark solution's appeal is due to the fact that both  $SW^b$  and  $\pi^b$  approximate  $SW^*$  up to a  $\Theta(\frac{1}{1-\alpha})$  factor, which we can be shown using similar arguments as in the unit-demand case. This is still a non-trivial claim; even with personalized payments, it is not clear if we can ever approximate the optimum social welfare via profit. Now that we have a benchmark solution which achieves good welfare and revenue, but is not implementable using item pricing, our goal in this proof becomes to compute item prices which approximate this benchmark solution in both revenue and welfare.

*Step 2: Augmented Walrasian Equilibrium.* Our goal will be to use  $(\vec{x}^b, \vec{y}^b)$  as a guide to design a sequence of (item) pricing solutions that together behave like the benchmark solution, and return the 'best approximation' among these solutions. Towards this end, we introduce the notion of an *Augmented Walrasian Equilibrium* with dummy price  $p^d$ , which is a social welfare maximizing solution for an augmented problem, consisting of the original instance plus  $|T|$  dummy buyers having constant valuation  $p^d$ , one for each good  $t \in T$ . Our notion of an augmented Walrasian equilibrium can be viewed an extension of the Walrasian Equilibrium with reserve prices concept introduced in [21] to settings where buyers purchase bundles of arbitrary sizes. Unlike the unit-demand case where there is a linear relationship between the reserve price

and the revenue, bounding the profit in equilibrium solutions (once dummy buyers are removed) with multi-minded buyers and production costs require a careful characterization of the various goods. Specifically, there are really two types of goods in such solutions: ones that are *saturated*, i.e.,  $p_t = c_t(y_t)$ , and ones which are artificially limited by the dummy price  $p^d$ , so  $p_t = p^d$ .

*Step 3: Starting dummy price.* The plan is to form a sequence of solutions resulting from Augmented Walrasian Equilibria with dummy prices, but what dummy price should we start with? Here we prove the following crucial claim. The main argument which makes this claim work is the ability to charge any deviation from the benchmark solution to the saturated goods in the current pricing solution, which are guaranteed to provide good welfare.

**Claim 5.2.** *Suppose that  $(\vec{p}, \vec{x}, \vec{y})$  denotes a pricing solution for the original instance obtained via an augmented Walrasian equilibrium with dummy price  $\frac{\vec{p}}{2\ell^{max}}$ . Then,*

$$SW^* - SW(\vec{p}, \vec{x}, \vec{y}) \leq \left(5 + \frac{6}{1 - \alpha}\right) [\pi(\vec{p}, \vec{x}, \vec{y}) + \pi(\vec{p}^*, \vec{x}^*, \vec{y}^*)].$$

*Step 4: Sequence of Solutions.* Previously, we identified a suitable dummy price and proved that the solution obtained via the augmented equilibrium at this dummy price (to some extent) captures most of the welfare of the benchmark solution. What about the profit of this solution? When all of the bundles desired by buyers are of equal cardinality, one can immediately show that the solution's profit also mimics that of the personalized payment scheme in the benchmark solution,  $\pi^b$ . When there is large disparity in the bundle sizes, however, the dummy price may result in buyers gravitating towards the smaller sized bundles. How do we fix this?

We next consider augmented Walrasian equilibria at dummy prices that are scaled versions of the original dummy price  $\frac{\vec{p}}{2\ell^{max}}$ , and once again, charge the lost welfare (due to a buyer facing high prices compared to her personalized payment in the benchmark solution) special class of saturated goods, which always yield high revenue. The profit due to the saturated goods is no larger than the social welfare of the solution, and thus, in either event, the social welfare cannot be small. More formally, we define a series of pricing solutions  $(\vec{p}(j), \vec{x}(j), \vec{y}(j))$  for  $j = 1$  to  $j = \delta = \lceil \log(\Delta) \rceil + 1$  to be the pricing solution for the original instance obtained via the augmented Walrasian equilibrium at dummy price  $2^j \frac{\vec{p}}{2\ell^{max}}$ . Let us also define  $SW(j)$  and  $\pi(j)$  to be the social welfare and profit of the respective pricing solutions. Then, we have that:

**Claim 5.3.** *For every  $j \in [0, \delta]$ ,  $SW(j) - SW(j + 1) \leq 3\pi(j) + 3\pi(j + 1)$ .*

Since we can bound the difference in welfare via profit, we can sequentially build solutions with good social welfare until we either reach a solution with good profit or  $j = 1 + \delta$  is reached, and we have no more solutions to construct. Now combining the above claims along with an upper bound on  $SW(\delta + 1)$  in terms of  $\pi(\delta + 1)$  (which we prove in the Appendix), we obtain a bound on  $SW^*$  in terms of the profits of various solutions, which would imply that the max-profit solution approximates the optimum welfare up to the desired bound.

**Claim 5.4.**

$$SW^* \leq \left(8 + 2\left(\frac{1}{1 - \alpha}\right)^{\frac{1}{\alpha}} + 4\frac{1}{1 - \alpha}\right) \left[\sum_{j=0}^{1+\delta} \pi(j) + \pi(\vec{p}^*, \vec{x}^*, \vec{y}^*)\right].$$

*Step 5: Final Pricing Algorithm.* Unfortunately, the above argument is still not strong enough to give us a bicriteria bound where the welfare is independent of  $\Delta$ . In the statement of Claim 5.4,

there are  $2 + \log(\Delta)$  profit terms on the right hand side; therefore, the solution giving maximum profit among  $\pi(0), \pi(1), \dots, \pi(\delta + 1)$  and  $\pi(\vec{p}^*, \vec{x}^*, \vec{y}^*)$  must yield a  $(2 + \log(\Delta))\Theta(\frac{1}{1-\alpha})$  approximation to both profit and welfare.

In our final step of the proof, we take this proof further and show that if the max-profit solution from the above claim does not result in a good welfare, then one can in fact identify another such solution with similar profit guarantees but *much better* welfare. Specifically, we show that by using some careful analysis, one can instead compute a solution whose approximation for social welfare is simply  $\Theta(\frac{1}{1-\alpha})$  and thus, is independent of  $\Delta$ . We now define our main algorithm: for the purpose of continuity, let  $\pi(-1)$  represent  $\pi(\vec{p}^*, \vec{x}^*, \vec{y}^*)$ .

1. Let  $k$  be the smallest index in the range  $[-1, 1 + \delta]$  such that  $\frac{SW^*}{\pi(k)}$  is no larger than  $2(\log(\Delta) + 2)\{8 + 2(\frac{1}{1-\alpha})^{\frac{1}{\alpha}} + 4\frac{1}{1-\alpha}\}$ .
2. Return the solution  $(\vec{p}(k), \vec{x}(k), \vec{y}(k))$ .

From Claim 5.4, it is clear that there exists at least one index  $k$  providing the desired guarantee. In fact, if  $k = -1$ , then we are done because the solution returned is the one maximizing social welfare. In our final claim (see Appendix), we show that for  $k$  defined as above, we have that  $SW^*$  is within a factor of  $12(\frac{2-\alpha}{1-\alpha})$  of  $SW(k)$ .  $\square$

Applying structural Claim 3.6 to the above proof, we get profit-welfare trade-offs analogous to Corollary 3.7, namely that for every given instance, there exists a constant  $1 \leq c \leq 12(\frac{2-\alpha}{1-\alpha})$ , so that if  $SW^* = cSW(Alg)$ , then  $SW^* \leq \frac{2c}{c-2}(5 + \frac{6}{1-\alpha})\pi(Alg)$ , when  $c \geq 2$ . For instance, this implies that either  $SW(Alg) \geq \frac{SW^*}{3}$  or that  $\pi(Alg)$  is actually a  $\Theta(\frac{1}{1-\alpha})$ -approximation to OPT. Here  $Alg$  refers to our bicriteria algorithm from Theorem 5.1.

### Consequences for other high-revenue solutions.

A direct application of Theorem 4.1 using our newly obtained bounds on multi-minded buyers as an intermediate yields the following claim.

**Claim 5.5.** *Let  $Alg$  be any algorithm that obtains at least as much profit as guaranteed by the algorithm of Theorem 5.1 on all instances. Then the social welfare obtained by  $Alg$  is at most a factor  $\Theta(\frac{\log(\Delta)}{1-\alpha})$  away from the optimum welfare.*

## 6 Conclusion and Future Directions

In this work, we were able to provide envy-free posted pricing algorithms that simultaneously approximate both profit and social welfare for markets with quite general buyer valuations and production costs. Such multi-objective algorithms are extremely well-motivated in a variety of realistic settings, where revenue and welfare are closely interconnected. We used our profit-maximization guarantees as a black-box and showed that any solution with reasonable profit guarantees (including the maximum profit solution) generates good welfare. In the process, we provide a partial characterization of the exact friction between these two objectives.

## References

- [1] Vineet Abhishek and Bruce E. Hajek. Efficiency loss in revenue optimal auctions. In *Proceedings of the 49th IEEE Conference on Decision and Control, CDC 2010, December 15-17, 2010, Atlanta, Georgia, USA*, pages 1082–1087, 2010.
- [2] Gagan Aggarwal, Gagan Goel, and Aranyak Mehta. Efficiency of (revenue-)optimal mechanisms. In *Proceedings 10th ACM Conference on Electronic Commerce (EC-2009), Stanford, California, USA, July 6-10, 2009*, pages 235–242, 2009.
- [3] Elliot Anshelevich, Koushik Kar, and Shreyas Sekar. Envy-free pricing in large markets: Approximating revenue and welfare. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, pages 52–64, 2015.
- [4] Yoram Bachrach, Sofia Ceppi, Ian A. Kash, Peter Key, and David Kurokawa. Optimising trade-offs among stakeholders in ad auctions. In *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*, pages 75–92, 2014.
- [5] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Item pricing for revenue maximization. In *Proceedings 9th ACM Conference on Electronic Commerce (EC-2008), Chicago, IL, USA, June 8-12, 2008*, pages 50–59, 2008.
- [6] Saptarshi Bhattacharya, Koushik Kar, Joe H Chow, and Aparna Gupta. Extended second price auctions for plug-in electric vehicle (pev) charging in smart distribution grids. In *American Control Conference (ACC), 2014*, pages 908–913. IEEE, 2014.
- [7] Avrim Blum, Anupam Gupta, Yishay Mansour, and Ankit Sharma. Welfare and profit maximization with production costs. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 77–86, 2011.
- [8] Patrick Briest. Uniform budgets and the envy-free pricing problem. In *ICALP 2008*.
- [9] Patrick Briest and Piotr Krysta. Buying cheap is expensive: Approximability of combinatorial pricing problems. *SIAM J. Comput.*, 40(6):1554–1586, 2011.
- [10] Niv Buchbinder and Rica Gonen. Incentive compatible multi-unit combinatorial auctions: A primal dual approach. *Algorithmica*, 72(1):167–190, 2015.
- [11] Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *The American Economic Review*, 86(1):180–194, 1996.
- [12] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of EC, 2007*.
- [13] Ning Chen, Xiaotie Deng, Paul W Goldberg, and Jinshan Zhang. On revenue maximization with sharp multi-unit demands. *Journal of Combinatorial Optimization*, pages 1–32, 2014.
- [14] Maurice Cheung and Chaitanya Swamy. Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. In *Proceedings of FOCS'08*, pages 35–44. IEEE, 2008.

- [15] Richard Cole and Shramas Rao. Applications of alpha-strongly regular distributions to bayesian auctions. In *Web and Internet Economics - 11th International Conference, WINE 2015, Amsterdam, The Netherlands, December 9-12, 2015, Proceedings*, pages 244–257, 2015.
- [16] Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 243–252, 2014.
- [17] Constantinos Daskalakis and George Pierrakos. Simple, optimal and efficient auctions. In *Internet and Network Economics - 7th International Workshop, WINE 2011, Singapore, December 11-14, 2011. Proceedings*, pages 109–121, 2011.
- [18] Xiaotie Deng, Paul W. Goldberg, Yang Sun, Bo Tang, and Jinshan Zhang. Pricing ad slots with consecutive multi-unit demand. In *Algorithmic Game Theory - 6th International Symposium, SAGT 2013, Aachen, Germany, October 21-23, 2013. Proceedings*, pages 255–266, 2013.
- [19] Ilias Diakonikolas, Christos H. Papadimitriou, George Pierrakos, and Yaron Singer. Efficiency-revenue trade-offs in auctions. In *Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part II*, pages 488–499, 2012.
- [20] Shaddin Dughmi, Tim Roughgarden, and Mukund Sundararajan. Revenue submodularity. *Theory of Computing*, 8(1):95–119, 2012.
- [21] Venkatesan Guruswami, Jason D Hartline, Anna R Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *SODA 2005*.
- [22] Jason D. Hartline and Vladlen Koltun. Near-optimal pricing in near-linear time. In *Algorithms and Data Structures, 9th International Workshop, WADS 2005, Waterloo, Canada, August 15-17, 2005, Proceedings*, pages 422–431, 2005.
- [23] Justin Hsu, Jamie Morgenstern, Ryan M. Rogers, Aaron Roth, and Rakesh Vohra. Do prices coordinate markets? In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 440–453, 2016.
- [24] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 45–60, 2015.
- [25] Sungjin Im, Pinyan Lu, and Yayun Wang. Envy-free pricing with general supply constraints for unit demand consumers. *J. Comput. Sci. Technol.*, 27(4):702–709, 2012.
- [26] Robert Kleinberg and Yang Yuan. On the ratio of revenue to welfare in single-parameter mechanism design. In *ACM Conference on Electronic Commerce, EC '13, Philadelphia, PA, USA, June 16-20, 2013*, pages 589–602, 2013.
- [27] Piotr Krysta, Orestis Telelis, and Carmine Ventre. Mechanisms for multi-unit combinatorial auctions with a few distinct goods. *J. Artif. Intell. Res. (JAIR)*, 53:721–744, 2015.



- [28] Anton Likhodedov and Tuomas Sandholm. Mechanism for optimally trading off revenue and efficiency in multi-unit auctions. In *Proceedings 5th ACM Conference on Electronic Commerce (EC-2004)*, New York, NY, USA, May 17-20, 2004, pages 268–269, 2004.
- [29] Brendan Lucier, Renato Paes Leme, and Éva Tardos. On revenue in the generalized second price auction. In *Proceedings of WWW 2012*.
- [30] Tim Roughgarden and Mukund Sundararajan. Is efficiency expensive. In *Third Workshop on Sponsored Search Auctions*, 2007.
- [31] Balasubramanian Sivan, Vasilis Syrgkanis, and Omer Tamuz. Lower bounds on revenue of approximately optimal auctions. In *Internet and Network Economics - 8th International Workshop, WINE 2012, Liverpool, UK, December 10-12, 2012. Proceedings*, pages 526–531, 2012.
- [32] Wayes Tushar, Walid Saad, H Vincent Poor, and David B Smith. Economics of electric vehicle charging: A game theoretic approach. *Smart Grid, IEEE Transactions on*, 3(4):1767–1778, 2012.

## A Appendix: Proof of Theorem 3.2 for Unit-Demand Buyers

We begin by defining some notation pertinent to the proof. Given a price vector  $\vec{p}$ , we will use  $q_i(\vec{p})$  to denote the price of the minimally priced good desired by buyer type  $i$ , i.e.,  $q_i(\vec{p}) = \min_{t \in B_i} p_t$ . For the rest of the proof, we will use the terms buyer and buyer type interchangeably. Recall that under any solution  $(\vec{p}, \vec{x}, \vec{y})$ , if buyer  $i$  purchases non-zero amount of good  $t \in B_i$ , then  $p_t = q_i(\vec{p}) = \lambda_i(x_i)$ . More specifically, we will only consider solutions that minimize the overall cost subject to the constraint that each buyer  $i$  only purchases from the minimally priced goods in  $S_i$ . That is, given  $\vec{p}, \vec{x}$ , there are several candidate allocation vectors  $\vec{y}$  that correspond to valid distributions of the demand vector  $\vec{x}$  on the goods. Unless mentioned otherwise, all of our solutions will consider allocation vectors that minimize the total cost  $C(\vec{y}) = \sum_{t \in T} C_t(y_t)$  among the set of valid allocation vectors. Observe that in such a solution, if buyer  $i$  purchases two different goods  $t, t'$ , then it has to be the case that the marginal cost of the two goods are equal, i.e.,  $c_t(y_t) = c_{t'}(y_{t'})$ . Therefore, without loss of generality, we can use  $r_i(\vec{y})$  to denote the marginal cost of any of the goods being used by buyer  $i$  under the given allocation. We will also use  $C(\vec{y})$  as a shortcut for the total cost incurred under a given allocation  $\vec{y}$ , i.e.,  $C(\vec{y}) = \sum_{t \in T} C_t(y_t)$ .

### The Social Welfare Maximizing Solution

The optimum welfare solution  $OPTW := (\vec{p}^*, \vec{x}^*, \vec{y}^*)$  is the pricing solution that maximizes the total social welfare of the system. The following simple claim (inspired by a similar claim from [3]) provides a useful characterization of the optimum solution.

**Claim A.1.** *There exists an optimum solution  $(\vec{p}^*, \vec{x}^*, \vec{y}^*)$  where for every good  $t \in T$ ,  $p_t^* = c_t(y_t^*)$ .*

We also define the notion of a welfare-maximizing solution constrained by demand. Specifically, an allocation vector  $\vec{y}$  is said to be a welfare-maximizer with respect to given demand  $\vec{x}$ , if  $\vec{y}$  is the minimum cost allocation consistent with the demand vector  $\vec{x}$ .

### Computing the Welfare Maximizing Solution

The algorithm described in Section 3 explicitly uses the prices in the social welfare maximizing solution  $OPTW$ . In order to compute the prices efficiently, we use the following convex program to compute the welfare maximizing demand and allocation vector  $(\vec{x}^*, \vec{y}^*)$  and then apply Claim A.1 to obtain the corresponding price vector.

$$\begin{aligned}
 & \max \sum_{i \in B} \int_{x=0}^{x_i} \lambda_i(x) dx - \sum_{t \in T} C_t(y_t) \\
 & \text{s.t.} \quad \sum_{t: t \in B_i} x_{it} = x_i \quad \forall i \in B \\
 & \quad \quad \sum_{i: i \in B_t} x_{it} = y_t \quad \forall t \in T \\
 & \quad \quad x_{it}, x_i, y_t \geq 0, \quad \forall i \in B, t \in T
 \end{aligned} \tag{1}$$

We now proceed with the proof of the theorem. The proof conforms to the following structure: first, we define a special pricing strategy known as a thresholded pricing vector and prove

that such a strategy satisfies some desirable properties. After this, we impose some sufficient conditions on such a pricing vector that allow us to prove bicriteria approximation bounds on the revenue and welfare. Finally, we prove that the thresholded pricing strategy presented in Section 3 does indeed satisfy these sufficient conditions.

*Proof.* The main instrument that allows us to design a simple but efficient pricing strategy is the notion of a *thresholded pricing vector*. Such a pricing vector extends uniform pricing strategies (all goods have the same price), which have been successfully applied to design solutions with good revenue [5] towards settings with production costs. We will show that thresholded pricing vectors enjoy many desirable properties and are convenient to analyze as opposed to the revenue maximizing solution, which is hard to get a grip on.

**Definition** (*Thresholded Pricing Vector*) A pricing vector  $\vec{p}$  is said to be a thresholded pricing vector if  $\exists$  a primary price  $\tilde{p}$  such that for every good  $t$ ,  $p_t = \max(\tilde{p}, p_t^*)$ .

We begin with a useful lemma that allows us to characterize the allocation resulting from a thresholded pricing vector.

**Lemma A.2** (Lemma 3.3 from Section 3). *Suppose that  $((\tilde{p})_{t \in T}, (\tilde{x})_{i \in B}, (\tilde{y})_{t \in T})$  is a pricing solution resulting from a thresholded pricing vector. Then,*

1. *The market can be clustered into two mutually disjoint sets of buyers and goods  $(B^H, T^H)$  and  $(B^L, T^L)$  so that the buyers in each cluster only purchase the goods in the same cluster and (a) for  $(i, t) \in (B^H, T^H)$ ,  $\tilde{p}_t = p_t^*$ ,  $\tilde{x}_i = x_i^*$ , and  $\tilde{y}_t = y_t^*$ ; (b) for  $(i, t) \in (B^L, T^L)$ ,  $\tilde{p}_t \geq p_t^*$ ,  $\tilde{x}_i \leq x_i^*$ , and  $c_t(\tilde{y}_t) \leq c_t(y_t^*)$ .*
2.  *$\vec{\tilde{y}}$  is a welfare-maximizing allocation vector with respect to the demand vector  $\vec{\tilde{x}}$ .*

*Proof.* Divide the buyers and goods as follows: let  $T^H$  be the set of goods with price (strictly) higher than  $\tilde{p}$  and let  $B^H$  be the buyers using these goods. Define  $T^L$  as the goods with price  $\tilde{p}$  and  $B^L$  as the corresponding buyers. Since each buyer only purchases from the min-priced goods available to her, buyers in  $B^L$  will only purchase (some of) the goods in  $T^L$ , and the same is true for those in  $B^H$ .

We begin by characterizing the solution corresponding to  $(B^H, T^H)$ : for all  $t \in T^H$ ,  $\tilde{p}_t = p_t^*$  as per the definition of  $\vec{\tilde{p}}$ . Consider any buyer  $i \in B^H$ ; clearly no good in  $T^L$  can belong to this buyer's desired set  $B_i$  since she is using the minimally priced good(s) available to her. This price must be exactly  $q_i(\vec{p}^*)$  and thus her demand must also be  $\tilde{x}_i = x_i^*$ . Next, for every good  $t' \in T^L$  and  $t \in T^H$ ,  $p_{t'}^* \leq \tilde{p} < p_t^*$ . Therefore, every buyer who is using a good in  $T^H$  in optimum welfare solution will still be using that good in our proxy solution. We can, therefore conclude that for all  $t \in T^H$ ,  $\tilde{y}_t = y_t^*$ .

Finally, look at any good  $t \in T^L$ , which is priced at exactly  $\tilde{p}$ . By definition,  $\tilde{p} \geq p_t^*$ . Therefore, for every buyer  $i \in T^L$ ,  $q_i(\vec{p}^*) \leq \tilde{p} = q_i(\vec{\tilde{p}})$ . And so,  $\tilde{x}_i \leq x_i^*$ . Since every buyer's demand is smaller in our solution w.r.t the max-welfare solution, we can use Lemma C.1 just for the buyers and goods in  $(B^L, T^L)$ . We get that for every  $t \in T^L$ ,  $c_t(\tilde{y}_t) \leq c_t(y_t^*)$ .

**(Part 2):** In order to show that  $\vec{\tilde{y}}$  has the smallest cost among all feasible allocations satisfying demand vector  $\vec{\tilde{x}}$ , it is sufficient (and necessary) to prove that for every buyer  $i$ , and every good  $t$  such that  $t \in B_i$ ,  $r_i(\vec{\tilde{y}}) \leq c_t(\tilde{y}_t)$ . Recall that every good  $t$  is priced at either  $p_t^*$  or at  $\tilde{p}$  if  $\tilde{p} > p_t^*$ .

By definition, we have two min-cost sub-allocations: 1) buyers in  $B^L$  are using the cheapest possible allocation using only the items in  $T^L$  because all goods in  $T^L$  have the same price

and we consider only cost minimizing allocations for the given pricing solution; 2) the same is true for  $B^H$  and  $T^H$ , this is because for these entities, the allocation is the same as in the max-welfare solution, which must definitely minimize cost for any given subset of the actual buyers and goods. So, in order to prove that this is a min-cost flow, we only need to consider the cross-edges, i.e., buyers belonging to  $B^L$  and goods desired by these buyers belonging to  $T^H$ . We already know that for any given buyer in  $B^H$  and good  $t \in T^L$ ,  $t \notin B_i$ . So we can conveniently ignore this case. What about the reverse case, can there be a buyer  $i$  in  $B^L$  and an item  $t$  in  $T^H$  such that  $c_t(\tilde{y}_t) = c_t(y_t^*) < r_i(\vec{y})$ ?

We know that for every  $t' \in T^L$ ,  $c_{t'}(\tilde{y}_{t'}) \leq c_{t'}(y_{t'}^*)$ . Since all buyers in  $B^L$  are using these goods in both the current solution and in  $(\vec{p}^*, \vec{x}^*, \vec{y}^*)$ , it must be that for every such buyer  $i \in B^L$ ,  $r_i(\vec{y}) \leq r_i(y^*)$ . But since  $i$  does not use any good  $t \in T^H$  even in the optimum solution, we get that  $r_i(y^*) \leq c_t(y_t^*) = c_t(\tilde{y}_t)$ . This completes the proof.  $\square$

Now that we have a better understanding of thresholded pricing solutions, i.e., solutions resulting from thresholded pricing vectors, we are ready to prove our main theorem ala Lemma 3.4 by constructing a suitable threshold vector and establishing upper bounds on  $SW(Alg)$  and  $SW^* - SW(Alg)$ . We divide the proof into three major claims: in the first two claims, we consider a general thresholded pricing solution and establish sufficient conditions that allow us to derive useful upper bounds on  $\pi(Alg)$  as a function of the primary price  $\tilde{p}$ . Following this, we construct an actual thresholded pricing vector that satisfies the required sufficient conditions and complete the proof.

**Claim A.3.** *Suppose that  $(\vec{p}, \vec{x}, \vec{y})$  is a thresholded pricing solution with primary price  $\tilde{p}$ . Then, the total social welfare of this solution  $SW(\vec{p}, \vec{x}, \vec{y})$  is at most a factor  $2^{\frac{\lambda^{max}}{\tilde{p}}} - 1$  times the profit due to this solution  $\pi(\vec{p}, \vec{x}, \vec{y})$ .*

*Proof.* The social welfare of the current solution is  $\sum_{i \in B} u_i(\tilde{x}_i) - C(\vec{y})$ . The function  $u_i$  is concave for every  $i \in B$ , and therefore  $u_i(\tilde{x}_i) \leq u'_i(0) \cdot \tilde{x}_i = \lambda^{max} \cdot \tilde{x}_i$ . So, our first inequality is the following,

$$\sum_{i \in B} u_i(\tilde{x}_i) - C(\vec{y}) \leq \sum_{i \in B} \lambda^{max} \tilde{x}_i - C(\vec{y}) = \frac{\lambda^{max}}{\tilde{p}} \sum_{i \in B} \tilde{p} \tilde{x}_i - C(\vec{y}) \leq \frac{\lambda^{max}}{\tilde{p}} \sum_{t \in T} \tilde{p} \tilde{y}_t - C(\vec{y}).$$

The final inequality comes from the fact that  $\tilde{p} \leq \tilde{p}_t$  for all  $t \in T$  and from rearranging the allocation from the buyers to the goods. Now, the total profit that the seller makes at the given strategy  $\pi(\vec{p}, \vec{x}, \vec{y})$  equals  $\sum_{t \in T} \tilde{p}_t \tilde{y}_t - C(\vec{y})$ . Using this, we get the following upper bound for the ratio of the welfare to profit

$$\begin{aligned} \frac{\sum_{i \in B} u_i(\tilde{x}_i) - C(\vec{y})}{\pi(\vec{p}, \vec{x}, \vec{y})} &\leq \frac{\frac{\lambda^{max}}{\tilde{p}} \sum_{t \in T} \tilde{p}_t \tilde{y}_t - C(\vec{y})}{\sum_{t \in T} \tilde{p}_t \tilde{y}_t - C(\vec{y})} \\ &\leq \frac{\frac{\lambda^{max}}{\tilde{p}} \sum_{t \in T} \tilde{p}_t \tilde{y}_t - \sum_{t \in T} \frac{1}{2} c_t(\tilde{y}_t) \tilde{y}_t}{\sum_{t \in T} [\tilde{p}_t - \frac{1}{2} c_t(\tilde{y}_t)] \tilde{y}_t} \\ &\leq \frac{\sum_{t \in T} \frac{\lambda^{max}}{\tilde{p}} \tilde{p}_t \tilde{y}_t - \frac{1}{2} \tilde{p}_t \tilde{y}_t}{\sum_{t \in T} [\tilde{p}_t \tilde{y}_t - \frac{1}{2} \tilde{p}_t \tilde{y}_t]} \\ &= 2 \frac{\lambda^{max}}{\tilde{p}} - 1. \end{aligned}$$

The second inequality above comes from the definition of doubly convex cost functions according to which  $C_t(\tilde{y}_t) \leq \frac{1}{2}c_t(\tilde{y}_t)$ . The third inequality comes from the fact that for every  $t$ ,  $c_t(\tilde{y}_t) \leq \tilde{p}_t$ . To see why this is true, observe from Lemma 3.3 that for every  $t \in T$ ,  $c_t(\tilde{y}_t) \leq c_t(y_t^*) = p_t^* \leq \tilde{p}_t$ . This completes the proof of the first of our three claims.  $\square$

**Claim A.4.** Suppose that  $(\vec{p}, \vec{x}, \vec{y})$  is a thresholded pricing solution with primary price  $\tilde{p}$  satisfying the following condition,

- For every  $i \in B$ , either  $\frac{\lambda_i(\tilde{x}_i) - r_i(\vec{y})}{|\lambda'_i(\tilde{x}_i)|} \leq \tilde{x}_i$  or  $\lambda_i(x_i^*) \geq \tilde{p}$ .

Then, the difference in social welfare with respect to  $OPTW$ , i.e.,  $SW_2 := SW^* - SW(\vec{p}, \vec{x}, \vec{y})$  is at most a factor  $\frac{1}{1-\alpha}$  times the profit due to this solution  $\pi(\vec{p}, \vec{x}, \vec{y})$ .

*Proof.* We need to prove an upper bound on  $SW^* - SW(\vec{p}, \vec{x}, \vec{y}) = \sum_{i \in B} (u_i(x_i^*) - u_i(\tilde{x}_i)) - (C(y^*) - C(\vec{y}))$ . Applying Lemma C.2, we get that  $(C(y^*) - C(\vec{y})) \geq \sum_{i \in B} r_i(\vec{y})(x_i^* - \tilde{x}_i)$ . Writing the utility as an integral, the required welfare can be simplified as

$$SW_2 \leq \sum_{i \in B} \left[ \int_{\tilde{x}_i}^{x_i^*} \lambda_i(x) dx - r_i(\vec{y})(x_i^* - \tilde{x}_i) \right] = \sum_{i \in B} \int_{\tilde{x}_i}^{x_i^*} (\lambda_i(x) - r_i(\vec{y})) dx.$$

Now, consider the function  $f_i(x) = \lambda_i(x) - r_i(\vec{y})$ . The second term is constant w.r.t  $x$  and  $\lambda_i$  is  $\alpha$ -SR and so from Lemma D.1,  $f_i$  is also  $\alpha$ -SR for all  $i \in B$ . Next, we can use our crucial Lemma D.2 to bound the integral as follows  $\int_{\tilde{x}_i}^{x_i^*} f_i(x) dx \leq \frac{1}{1-\alpha} \left( \frac{f_i(\tilde{x}_i)}{|f'_i(\tilde{x}_i)|} \right) (f_i(\tilde{x}_i) - f_i(x_i^*))$ . Summing this up, one obtains,

$$SW_2 \leq \frac{1}{1-\alpha} \sum_{i \in B} \frac{\lambda_i(\tilde{x}_i) - r_i(\vec{y})}{|\lambda'_i(\tilde{x}_i)|} (\lambda_i(\tilde{x}_i) - \lambda_i(x_i^*)).$$

Let  $B^L$  and  $B^H$  be as in Lemma 3.3. For  $i \in B^H$ , we know that  $\lambda_i(\tilde{x}_i) = \tilde{p}_i = p_i^* = \lambda_i(x_i^*)$ , so the above expression equals 0. From the claim statement, we know that for  $i \in B^L$ , the quantity  $\frac{\lambda_i(\tilde{x}_i) - r_i(\vec{y})}{|\lambda'_i(\tilde{x}_i)|}$  must be bounded from above by  $\tilde{x}_i$ . Thus, we have that,

$$SW_2 \leq \frac{1}{1-\alpha} \sum_{i \in B^L} \tilde{x}_i (\lambda_i(\tilde{x}_i) - r_i(\vec{y})).$$

Now, we can complete the proof using the fact that for all  $i$ ,  $\lambda_i(\tilde{x}_i)\tilde{x}_i$  is the total profit due to that buyer.

$$SW_2 \leq \frac{1}{1-\alpha} \sum_{t \in T} [\tilde{y}_t \tilde{p}_t - c_t(\tilde{y}_t) \tilde{y}_t] \leq \frac{1}{1-\alpha} \left[ \sum_{t \in T} \tilde{y}_t \tilde{p}_t - C(\vec{y}) \right] = \frac{1}{1-\alpha} \pi(\vec{p}, \vec{x}, \vec{y}).$$

$\square$

We have shown that if we can conjure up a nice thresholded price vector that satisfies certain properties, we can bound the social welfare of the optimum solution in terms of the profit obtained by the above solution. We now show that the price vector described in Section 3 satisfies all of our requirements.

**Claim A.5.** *We can compute in poly-time a solution  $((\tilde{p})_{t \in T}, (\tilde{x})_{i \in B}, (\tilde{y})_{t \in T})$  satisfying the following two conditions,*

1.  $\exists$  a price  $\tilde{p}$  such that for every  $t \in T$ ,  $\tilde{p}_t = \max(p_t^*, \tilde{p})$ .
2. For every  $i \in B$ , either  $\frac{\lambda_i(\tilde{x}_i) - r_i(\tilde{y})}{|\lambda'_i(\tilde{x}_i)|} \leq x_i$  or  $\lambda_i(x_i^*) \geq \tilde{p}$ .

*Proof.* Suppose that all the users have  $\alpha$ -strongly regular utilities. Then, set  $\tilde{p} = (1 - \alpha)^{\frac{1}{\alpha}} \lambda^{max}$ . Construct the price vector  $\vec{\tilde{p}}$  as described above. Moreover, let  $\vec{\tilde{x}}$  and  $\vec{\tilde{y}}$  be the corresponding buyer demand vector and (feasible) allocation vector at this price. All that is remaining is to prove that this solution satisfies the second condition mentioned in the claim.

Consider any  $i \in B$ : the min-priced good available to this buyer is priced at either  $\tilde{p}$  or  $q_i(\vec{p}^*)$ . In the  $\tilde{p}$  case, we know that the user has a demand of  $\tilde{x}_i$  units satisfying  $\lambda_i(\tilde{x}_i) = \tilde{p} = \lambda^{max}(1 - \alpha)^{\frac{1}{\alpha}}$ . By definition  $\lambda^{max} = \lambda_i(0)$  for all  $i$ . Therefore, for every user facing a price of  $\tilde{p}$  we have that

$$\lambda_i(0) = \left(\frac{1}{1 - \alpha}\right)^{1/\alpha} \lambda_i(\tilde{x}_i).$$

We can loosen the above inequality as follows:  $\lambda_i(0) - r_i(\vec{\tilde{y}}) \geq \left(\frac{1}{1 - \alpha}\right)^{1/\alpha} (\lambda_i(\tilde{x}_i) - r_i(\vec{\tilde{y}}))$ . Now take  $f_i(x) = \lambda_i(x) - r_i(\vec{\tilde{y}})$  and apply corollary D.5 with respect to  $\tilde{x}_i$ . We get that  $\frac{f_i(\tilde{x}_i)}{|f'_i(\tilde{x}_i)|} \leq \tilde{x}_i$ , which is the required criterion.

For the second case, the proof follows trivially because the min-priced good  $t$  available to the user  $i$  has a price of  $p_t^* > \tilde{p}$ . This is the same as the min-priced good available to this user in the optimum solution and therefore  $\lambda_i(\tilde{x}_i) = \lambda_i(x_i^*) = p_t^* \geq \tilde{p}$ . □

## Final Leg: Combining all of the Claims

The rest of the proof follows from a direct application of our structural lemmas. Primarily, consider the pricing algorithm  $Alg$ , that chooses a pricing vector as defined in Claim A.5. Then, from Claims A.3 and A.4, we get that  $SW(Alg) \leq \left((2\frac{1}{1-\alpha})^{\frac{1}{\alpha}} - 1\right) \pi(Alg)$  and  $SW^* - SW(Alg) \leq \frac{1}{1-\alpha} \pi(Alg)$ . Applying Lemma 3.4 yields the required bicriteria result.

Finally, we claim that  $\zeta = \Theta(\frac{1}{1-\alpha})$ . To see why, consider the function  $\frac{2(\frac{1}{1-\alpha})^{\frac{1}{\alpha}}}{\frac{1}{1-\alpha}}$ . Upon differentiation, we infer that this function is non-increasing and its maximum value of  $2e$  is obtained as  $x \rightarrow 0$ . □

## B Appendix: Proof of Theorem 5.1 for Multi-Minded Buyers

**Proof Sketch:** The proof for the general case with multi-minded buyers is quite involved and so we begin by providing a high level overview of the various ingredients that combine to form the proof. The first step involves defining a benchmark solution  $(\vec{x}^b, \vec{y}^b)$  whose social welfare, as in the unit-demand case, is at most a  $\frac{2-\alpha}{1-\alpha}$  factor away from  $SW^*$ . In the unit-demand case, we were able to identify a suitable price vector to actually implement this benchmark solution and extract a good fraction of its welfare as revenue. However, this may no longer be possible for the general case as an analogous pricing solution would involve personalized payments for each buyer

and thus may not be realizable using simple item pricing. In other words, the benchmark solution is only realizable by charging different prices to different buyers for the same good, instead of item pricing in which a good has a single price for all buyers. Charging such discriminatory prices would change the problem completely and give the seller much more power (for example, maximizing revenue becomes almost trivial, instead of NP-Complete). Our goal in this proof is therefore to compute item prices which approximate this benchmark solution in both revenue and welfare.

The rest of the proof does not really use the benchmark solution directly; instead our goal will be to use  $(\vec{x}^b, \vec{y}^b)$  as a guide to design a sequence of (item) pricing solutions that together behave like the benchmark solution, and return the ‘best approximation’ among these solutions. Towards this end, we introduce the notion of an *Augmented Walrasian Equilibrium* with dummy price  $p$ , which is a social welfare maximizing solution for an augmented problem, consisting of the original instance plus  $|T|$  dummy buyers having constant valuation  $p$ , one for each good  $t \in T$ . One way to think about this is as a Walrasian equilibrium with a reserve price  $p$  for each good, so that the prices are required to be above  $p$ . We show that these augmented equilibria have several desirable properties once the dummy buyers are removed from the final solution.

The starting point for our algorithm is the identification of a carefully chosen dummy price (see Claim 5.2) so that the ensuing augmented equilibrium (minus the dummy buyers) captures the welfare of the benchmark solution to a large extent. Of course, the solution may not extract the required profit, and to correct this (Claim 5.3), we show that simply scaling the dummy price by a factor of two yields a solution whose welfare still approximates the original solution. The repeated scaling gives us a sequence of  $\log(\Delta)$  different pricing solutions. Finally, we show that at least one solution in this sequence has profit and welfare which both approximate that of the benchmark solution.

Both of the above claims make use of a charging argument that may be of independent interest, wherein the lost welfare (due to a buyer facing high prices compared to her personalized payment in the benchmark solution) is charged to a special class of goods, referred to as saturated goods, which always yield high revenue.

### *Proof. Benchmark Solution*

Recall that  $(\vec{p}^*, \vec{x}^*, \vec{y}^*)$  denotes the pricing solution achieving the optimum social welfare. Define  $\tilde{p} = \lambda^{\max}(1 - \alpha)^{\frac{1}{\alpha}}$ . Now, the benchmark demand vector  $\vec{x}^b$  is defined as follows: for each buyer  $i$ ,  $x_i^b = \min\{x_i^*, \lambda^{-1}(\tilde{p})\}$ . The allocation vector  $\vec{y}^b$  is simply a scaled-down version of the optimum allocation vector, i.e., suppose that  $x_i^*(S)$  is amount of bundle  $S$  consumed by buyer type  $i$  in  $(\vec{p}^*, \vec{x}^*, \vec{y}^*)$ . Then, the amount of bundle  $S$  consumed by  $i$  in the benchmark solution  $x_i^b(S) := \frac{x_i^b}{x_i^*} x_i^*(S)$ . This consumption pattern gives rise to the allocation vector  $\vec{y}^b$  on the goods. Let  $SW^b$  denote the social welfare of this benchmark solution.

Does there exist a pricing solution  $\vec{p}^b$  with a single price per good that implements the above benchmark solution? Unfortunately, one can easily define instances where this is not the case. That said, a trivial way to implement the benchmark solution is via a personalized payment scheme, i.e., each buyer  $i$  is charged a price of  $\lambda_i(x_i^b)$  that is completely independent of other buyers. For convenience, we define  $\pi^b = \sum_{i \in B} \lambda_i(x_i^b) x_i^b - C(\vec{y}^b)$  to be the profit due to the personalized payments. We begin by highlighting an obvious property of the benchmark solution that will come in handy later on.

**Proposition B.1.** *The benchmark solution vectors are dominated by the optimum solution vectors, i.e.,  $\vec{x}^* \geq \vec{x}^b$  and  $\vec{y}^* \geq \vec{y}^b$ . Therefore for every good  $t$ ,  $c_t(y_t^b) \leq c_t(y_t^*)$ .*

Despite its lack of implementability, the benchmark solution's appeal lies in the fact that both  $SW^b$  and  $\pi^b$  approximate  $SW^*$  up to a  $\Theta(\frac{1}{1-\alpha})$  factor, which we formally state below. Observe that this is still a non-trivial claim; even with personalized payments, it is not clear if we can ever approximate the optimum social welfare via profit. For instance, for the max-welfare solution  $(\vec{x}^*, \vec{y}^*)$ , a personalized charging scheme achieves the same profit as  $\pi(\vec{p}^*, \vec{x}^*, \vec{y}^*)$ , which may be quite poor. Our goal however, will be to use  $(\vec{x}^b, \vec{y}^b)$  as a guide to design a sequence of (item) pricing solutions that together behave like the benchmark solution, and return the 'best approximation' among these solutions.

**Claim B.2.** (1)  $SW^b \leq \left[2\left(\frac{1}{1-\alpha}\right)^{\frac{1}{\alpha}} - 1\right] \pi^b$  (2)  $SW^* - SW^b \leq \frac{1}{1-\alpha} \pi^b$ .

The proof of the above claim is extremely similar to the proofs of Claims A.3 and A.4, and we do not explicitly prove it here as we do not really use the claim anywhere in the rest of this proof. Instead, we will state and prove approximate versions of the above claim that we will require later. These approximate versions shed additional light on some of the sufficient conditions that a pricing solution must satisfy in order to capture social welfare via profit.

**Claim B.3.** 1. Suppose that  $(\vec{p}, \vec{x}, \vec{y})$  is a pricing solution satisfying  $\lambda_i(x_i) \geq \tilde{p}$  for every  $i$ , and  $p_t \geq c_t(y_t)$  for all  $t \in T$ . Then,

$$SW(\vec{p}, \vec{x}, \vec{y}) \leq \left[2\left(\frac{1}{1-\alpha}\right)^{\frac{1}{\alpha}} - 1\right] \pi(\vec{p}, \vec{x}, \vec{y}).$$

2. Suppose that  $\vec{x}$  is a demand vector satisfying  $x_i \geq x_i^b$  for every  $i$ . Then,

$$\sum_{i \in B} \int_{x_i}^{x_i^*} \lambda_i(x) dx \leq \frac{1}{1-\alpha} \sum_{i \in B} \lambda_i(x_i) x_i.$$

*Proof. (Statement 1):*

The social welfare of  $(\vec{p}, \vec{x}, \vec{y})$  is  $\sum_{i \in B} u_i(x_i) - C(\vec{y})$ . The function  $u_i$  is concave for every  $i \in B$ , and therefore  $u_i(x_i) \leq u'_i(0) \cdot x_i = \lambda_i^{max} x_i$ . So, our first inequality is the following,

$$\sum_{i \in B} u_i(x_i) - C(\vec{y}) \leq \sum_{i \in B} \lambda_i^{max} x_i - C(\vec{y}) = \frac{\lambda^{max}}{\tilde{p}} \sum_{i \in B} \tilde{p} x_i - C(\vec{y}) = \left(\frac{1}{1-\alpha}\right)^{\frac{1}{\alpha}} \sum_{i \in B} \tilde{p} x_i - C(\vec{y}).$$

Now, the total profit that the seller makes at the given solution  $\pi(\vec{p}, \vec{x}, \vec{y})$  equals  $\sum_{i \in B} \lambda_i(x_i) x_i - C(\vec{y}) \geq \sum_{i \in B} \tilde{p} x_i - C(\vec{y})$ . Using this, we get the following upper bound for the ratio of the welfare to profit

$$\begin{aligned} \frac{\sum_{i \in B} u_i(x_i) - C(\vec{y})}{\pi(\vec{p}, \vec{x}, \vec{y})} &\leq \frac{\left(\frac{1}{1-\alpha}\right)^{\frac{1}{\alpha}} \sum_{i \in B} \lambda_i(x_i) x_i - C(\vec{y})}{\sum_{i \in B} \lambda_i(x_i) x_i - C(\vec{y})} \\ &\leq \frac{\left(\frac{1}{1-\alpha}\right)^{\frac{1}{\alpha}} \sum_{i \in B} \lambda_i(x_i) x_i - \frac{1}{2} \sum_{i \in B} \lambda_i(x_i) x_i}{\sum_{i \in B} \lambda_i(x_i) x_i - \frac{1}{2} \sum_{i \in B} \lambda_i(x_i) x_i} \\ &= 2\left(\frac{1}{1-\alpha}\right)^{\frac{1}{\alpha}} - 1. \end{aligned}$$



The second inequality above stems from Lemma C.3 in the Appendix, which states that for any feasible pricing solution,  $C(\vec{y}) \leq \frac{1}{2} \sum_{i \in B} \lambda_i(x_i)x_i$ . This completes the proof of the first statement.

**Statement 2:**

We need to prove an upper bound on the quantity  $\sum_{i \in B} \int_{x_i}^{x_i^*} \lambda_i(x)dx$  given the condition that  $x_i \geq x_i^b$  for all  $i \in B$ . Since  $\lambda_i$  is an  $\alpha$ -strongly regular function for all  $i$ , we can apply our crucial Lemma D.2 to bound the integral as follows  $\int_{x_i}^{x_i^*} \lambda_i(x)dx \leq \frac{1}{1-\alpha} \left( \frac{\lambda_i(x_i)}{|\lambda_i'(x_i)|} \right) (\lambda_i(x_i) - \lambda_i(x_i^*))$ . Summing this up, one obtains,

$$\begin{aligned} \sum_{i \in B} \int_{x_i}^{x_i^*} \lambda_i(x)dx &\leq \frac{1}{1-\alpha} \sum_{i \in B'} \frac{\lambda_i(x_i)}{|\lambda_i'(x_i)|} (\lambda_i(x_i) - \lambda_i(x_i^*)) \\ &\leq \frac{1}{1-\alpha} \sum_{i \in B'} x_i \lambda_i(x_i). \end{aligned}$$

In the above expression,  $B'$  is the set of buyers for whom  $\lambda_i(x_i) > \lambda_i(x_i^*)$ . For the remaining buyers, the integral is not positive and so we can safely ignore these terms. Further, since  $x_i \geq x_i^b$ , and  $\lambda_i(x_i) > \lambda_i(x_i^*)$  for all  $i \in B'$ , it must be that  $\lambda_i(x_i^b) > \lambda_i(x_i^*)$ , and so by definition of the benchmark solution, we have that  $\lambda(x_i^b) = \tilde{p}$ . Thus, it is not hard to see that for every buyer  $i \in B'$ ,  $\lambda_i(x_i) \leq \lambda_i(x_i^b) = \tilde{p} = \lambda^{max}(1-\alpha)^{\frac{1}{\alpha}}$ , and so, as per Corollary D.5,  $\lambda_i(x_i)/|\lambda_i'(x_i)| \leq x_i$ . This completes the proof.  $\square$

## Augmented Walrasian Equilibrium with Dummy Prices

The most important notion in this proof is that of an augmented Walrasian equilibrium. At a high level, the augmented equilibrium is a useful tool that allows us to approximate ‘personalized payment’-based solutions via item pricing as long as the total price charged to different buyers is correlated. For instance, in the benchmark solution, buyers are either charged  $\tilde{p}$  or  $\lambda_i(x_i^*)$ .

**Definition** Given an instance  $\mathcal{I} = (B, T)$ , an augmented Walrasian equilibrium with dummy price  $p$  is the social welfare maximizing allocation for an instance  $\mathcal{I}^+ = (B \cup B', T)$  where  $B, T$  are as defined previously and  $B'$  consists of  $|T|$  buyer types, one for each  $t \in T$ , such that the buyer  $i_t$  corresponding to good  $t$  has only good  $t$  in her desired set of bundles (exactly one bundle consisting of one item). Moreover, this buyer has a constant valuation  $\lambda_{i_t}(x) = p$  for  $x \leq L$ , where  $L$  is some sufficiently large population, and  $\lambda_{i_t}(x) = 0$  for  $x > L$ .

One can imagine that each new buyer type has a virtually infinite population of buyers all of whom desire one good and have the same valuation. Given  $\mathcal{I}, p$ , an augmented Walrasian equilibrium with this dummy price can be computed efficiently using the same convex program as before.

**Pricing Solutions Obtained via Augmented Equilibria** Given a dummy price  $p$ , the augmented equilibrium computation gives us a social welfare maximizing demand and allocation vector  $(\vec{x}^p, \vec{y}^p)$  for the corresponding instance  $\mathcal{I}^+$ . However, we are only interested in valid pricing solutions for our original instance  $\mathcal{I}$ . We now provide a simple transformation to convert  $(\vec{x}^p, \vec{y}^p)$  into a pricing solution  $(\vec{p}, \vec{x}, \vec{y})$  for the original instance, which we will take for granted in the rest of this work. Essentially, such a pricing solution is obtained by taking the augmented Walrasian equilibrium, pricing each good at its marginal cost, and then removing all consumption due to dummy buyers. Formally, define  $p_t = c_t(y_t^p)$  for all  $t \in T$ , and for  $i \in B$ ,  $x_i = x_i^p$ . Finally, let

these buyers consume the same bundles as in the augmented equilibrium, i.e., for every  $S \in B_i$ , the amount of this bundle consumed by the buyer in  $(\vec{p}, \vec{x}, \vec{y})$ ,  $x_i(S)$  is the same as the amount consumed by this buyer in the augmented equilibrium,  $x_i^p(S)$ .

It is evident from the definition that  $(\vec{p}, \vec{x}, \vec{y})$  is a valid pricing solution for our original instance  $\mathcal{I}$ . We first highlight some simple properties obeyed by such a pricing solution irrespective of the exact dummy price and in the process, gain some understanding of the structure of these solutions.

**Lemma B.4.** *Given an instance and a dummy price  $p^d$ , let  $(\vec{p}, \vec{x}, \vec{y})$  denote the pricing solution for the original instance obtained via the augmented Walrasian equilibrium with dummy price  $p^d$ . Then,*

1. *For every good  $t \in T$ , we have  $p_t = \max\{p^d, c_t(y_t)\}$ .*
2. *The profit due to the solution  $\pi(\vec{p}, \vec{x}, \vec{y}) \geq C(\vec{y})$ .*

*Proof.* Suppose that  $(x^{\vec{p}^d}, y^{\vec{p}^d})$  denotes the augmented Walrasian equilibrium with dummy price  $p^d$ . By construction, we know that for every  $t$ ,  $p_t = c_t(y_t^{p^d})$ .

1. We first claim that for every good  $t$ ,  $c_t(y_t^{p^d}) \geq p^d$ . Indeed, if this condition were not met for some  $t$ , then one could increase the social welfare of  $(x^{\vec{p}^d}, y^{\vec{p}^d})$  by simply allocating a non-zero amount of good  $t$  to the dummy buyer assigned to that good  $i_t$  violating the fact that  $(x^{\vec{p}^d}, y^{\vec{p}^d})$  is a social welfare maximizing allocation. Therefore for every  $t$ , it is true that  $p_t = c_t(y_t^{p^d}) \geq p^d$ .

Since buyers are using the same bundles as in the augmented equilibrium, the consumption of any good in  $\vec{y}$  cannot be larger than its consumption in  $y^{\vec{p}^d}$ , i.e., for all  $t \in T$ ,  $y_t \leq y_t^{p^d}$ . This implies that  $p_t = c_t(y_t^{p^d}) \geq c_t(y_t)$ .

Now consider some good  $t$  for which  $p_t > p^d$ . We claim that no dummy buyer can be consuming this good in the augmented equilibrium. Once again, we show this via contradiction. Assume to the contrary: then, one could increase the social welfare of  $(x^{\vec{p}^d}, y^{\vec{p}^d})$  by removing a non-zero amount of the allocation given to dummy buyer  $i_t$ , which violates the optimality of the solution. This means that the entire consumption of good  $t$  comes from the original buyers in  $B$ , and therefore,  $y_t = y_t^{p^d}$ . This in turn implies that  $c_t(y_t) = p_t$ .

2. We have already showed that for all  $t$ ,  $p_t \geq c_t(y_t)$ . Therefore, the second statement is a consequence of Lemma C.3.

□

**Saturated Goods** Perhaps, the most useful feature of pricing solutions obtained via augmented equilibria is highlighted in the second statement of the above lemma. Any good that is not priced at the dummy price is in a sense *saturated* as  $p_t = c_t(y_t)$ . We refer to such goods having  $p^d < p_t = c_t(y_t)$  as saturated goods. Saturated goods usually yield good profit (due to high consumption) and thus, are useful in bounding the social welfare in terms of profit. As we will show later, for a suitably chosen dummy price, any welfare lost from the benchmark solution can be charged to the income obtained only from these saturated goods.

## First shot at an actual dummy price

Now that we have a better understanding of pricing solutions obtained via augmented equilibria, we pose our first major question: can we effectively approximate the benchmark solution using a carefully chosen dummy price? Our journey towards answering the above question begins with the answer to a slightly relaxed question: what is the highest possible dummy price so that resulting solution has a social welfare comparable to that of the benchmark solution? Clearly, if the dummy price is small enough, one can recover the original optimum solution. The insistence on high prices is to ensure that each buyer's demand in the augmented equilibria is somewhat comparable to  $\lambda_i(x_i^b)$ . The following three claims shed some light on this question, by identifying  $\frac{\hat{p}}{2\ell^{max}}$  as the appropriate dummy price. At this dummy price, we show that any deviation from the benchmark solution can be charged to the revenue provided by the saturated goods (and so the deviation cannot be too large). Recall that  $\ell^{max}$  is the size of the largest bundle desired by any buyer type.

**Lemma B.5.** *Suppose that  $(\vec{x}^a, \vec{y}^a)$  is any given benchmark solution, satisfying  $\lambda_i(x_i^a) \geq \hat{p} > 0$  for all  $i \in B$ . Let  $(\vec{p}, \vec{x}, \vec{y})$  denote a pricing solution for the original instance obtained via an augmented Walrasian equilibrium with dummy price  $\frac{\hat{p}}{2\ell^{max}}$ . Then, for any buyer  $i$  having  $\lambda_i(x_i) > \lambda_i(x_i^a)$ , and every bundle  $S$  desired by  $i$  ( $S \in B_i$ ), we have that,*

- $S$  contains at least one saturated good.
- $\lambda_i(x_i) \leq 2 \sum_{t \in S \cap \hat{S}} p_t$ , where  $\hat{S}$  is the set of saturated goods in  $(\vec{p}, \vec{x}, \vec{y})$ .

*Proof.* The first part of the lemma follows from simple contradiction. Consider a buyer  $i$  and bundle  $S \in B_i$  as mentioned in the statement of the lemma, and suppose that  $S$  does not contain any saturated good. Therefore, as per Lemma B.4, it must be true that for every  $t \in S$ ,  $p_t = \frac{\hat{p}}{2\ell^{max}}$ . But since we have a valid pricing solution, it is also true that  $\lambda_i(x_i) \leq \sum_{t \in S} p_t$ . So, we have that

$$\hat{p} \leq \lambda_i(x_i^a) < \lambda_i(x_i) \leq \sum_{t \in S} p_t = |S| \frac{\hat{p}}{2\ell^{max}} \leq \frac{\hat{p}}{2},$$

which is a contradiction. Note that  $|S| \leq \ell^{max}$ , by definition.

Now that we have established that  $S$  contains at least one saturated good, we can write  $\lambda_i(x_i)$  as,

$$\lambda_i(x_i) \leq \sum_{t \in S \cap \hat{S}} p_t + \sum_{t \in S \setminus (S \cap \hat{S})} p_t \leq \sum_{t \in S \cap \hat{S}} p_t + \frac{\hat{p}}{2}.$$

Since  $\lambda_i(x_i) \geq \hat{p}$ , we get the desired lemma (second statement).  $\square$

Building on the previous lemma, we actually bound the welfare loss from the benchmark solution in terms of the income (profit without the cost) due to the saturated goods in  $(\vec{p}, \vec{x}, \vec{y})$  as well as  $(\vec{p}^*, \vec{x}^*, \vec{y}^*)$ . The introduction of *OPTW* here may appear strange at first but we remind the reader that bounding social welfare in terms of the profit at *OPTW* does not really hurt us by much as a high profit at *OPTW* leads to a trivial bicriteria approximation algorithm with optimal welfare.

**Lemma B.6.** *Suppose that  $(\vec{x}^a, \vec{y}^a)$  is any given benchmark solution and  $(\vec{p}, \vec{x}, \vec{y})$  denotes a pricing solution for the original instance obtained via an augmented Walrasian equilibrium with dummy price  $\frac{\hat{p}}{2\ell^{max}}$ . Also, suppose that the following conditions are satisfied*

1.  $\lambda_i(x_i^a) \geq \hat{p}$  for all  $i \in B$ .
2.  $\vec{x}^* \geq \vec{x}^a$  and  $\vec{y}^* \geq \vec{y}^a$ .

Let  $\hat{B}$  be the set of buyers for whom  $\lambda_i(x_i) > \lambda_i(x_i^a)$ . If  $\hat{S}$  denotes the set of saturated goods in  $(\vec{p}, \vec{x}, \vec{y})$ , then,

$$\sum_{i \in \hat{B}} \int_{x_i}^{x_i^a} \lambda_i(x) dx \leq \sum_{i \in \hat{B}} \lambda_i(x_i) x_i^a \leq 2 \sum_{t \in \hat{S}} p_t y_t + 2 \sum_{t \in \hat{S}} p_t^* y_t^*$$

*Proof.* Suppose that  $x_i^a(S)$  is the amount of bundle  $S$  consumed by buyer  $i$  in the benchmark solution. Then, we have that

$$\sum_{i \in \hat{B}} \lambda_i(x_i) x_i^a = \sum_{i \in \hat{B}} \sum_{S \in B_i} \lambda_i(x_i) x_i^a(S).$$

Moreover, we know that for every  $S \in B_i$ ,  $\lambda_i(x_i) \leq 2 \sum_{t \in S \cap \hat{S}} p_t$ , from Lemma B.5. Using this upper bound, we get

$$\sum_{i \in \hat{B}} \lambda_i(x_i) x_i^a \leq \sum_{i \in \hat{B}} \sum_{S \in B_i} \sum_{t \in S \cap \hat{S}} 2 p_t x_i^a(S) \leq 2 \sum_{t \in \hat{S}} p_t y_t^a.$$

Now, we divide the goods in  $\hat{S}$  into two categories: set  $\hat{S}^a$  represents the goods for which  $y_t \geq y_t^a$ , and the rest of the goods fall under  $\hat{S}^*$ . Moreover, for every good  $t \in \hat{S}^*$ ,  $c_t(y_t^*) \geq c_t(y_t^a) \geq c_t(y_t) = p_t$ . Using this, we can proceed with our final simplification.

$$2 \sum_{t \in \hat{S}} p_t y_t^a \leq 2 \sum_{t \in \hat{S}^a} p_t y_t^a + 2 \sum_{t \in \hat{S}^*} p_t y_t^a \leq 2 \sum_{t \in \hat{S}^a} p_t y_t + 2 \sum_{t \in \hat{S}^*} p_t^* y_t^*.$$

Since  $\hat{S}^*, \hat{S}^a \subseteq \hat{S}$ , we have our lemma.  $\square$

With our properties of augmented Walrasian equilibrium solutions firmly in place, we are now ready to prove our first crucial claim, a lower bound for the social welfare of the solution  $(\vec{p}, \vec{x}, \vec{y})$  obtained via an augmented equilibrium in terms of the optimum welfare as well as the profit at  $(\vec{p}, \vec{x}, \vec{y})$  as well as at  $OPTW$ . Assuming for the time being that the profit at  $OPTW$  is not that large, the claim below gives us a handle on a (high enough) dummy price for which our solution behaves like the benchmark solution. This dummy price will be the starting point, from which we design a sequence of solutions to extract good profit. If the current dummy price was much smaller (say), then we would still be guaranteed a good social welfare but unfortunately, good profit becomes an arduous task.

**Claim B.7** (Claim 5.2 from Section 5). *Suppose that  $(\vec{p}, \vec{x}, \vec{y})$  denotes a pricing solution for the original instance obtained via an augmented Walrasian equilibrium with dummy price  $\frac{\tilde{p}}{2\ell^{max}}$ . Then,*

$$SW^* - SW(\vec{p}, \vec{x}, \vec{y}) \leq \left(5 + \frac{6}{1 - \alpha}\right) [\pi(\vec{p}, \vec{x}, \vec{y}) + \pi(\vec{p}^*, \vec{x}^*, \vec{y}^*)].$$

Recall that  $\tilde{p} = \lambda^{max}(1 - \alpha)^{\frac{1}{\alpha}}$ .

*Proof.* Consider  $SW^* - SW(\vec{p}, \vec{x}, \vec{y})$ ; this can be expressed as follows,

$$SW^* - SW(\vec{p}, \vec{x}, \vec{y}) = \sum_{i \in B} \int_{x_i}^{x_i^*} \lambda_i(x) dx - (C(\vec{y}^*) - C(\vec{y})).$$

Recall our original benchmark solution  $(\vec{x}^b, \vec{y}^b)$ . Divide  $B$  into two subsets:  $\hat{B}$  denotes the set of buyers for whom  $\lambda_i(x_i) > \lambda_i(x_i^b)$  and  $B \setminus \hat{B}$  denotes the other buyers. We now further simplify the above equation,

$$\begin{aligned} \sum_{i \in B} \int_{x_i}^{x_i^*} \lambda_i(x) dx &= \sum_{i \in \hat{B}} \int_{x_i}^{x_i^*} \lambda_i(x) dx + \sum_{i \in B \setminus \hat{B}} \int_{x_i}^{x_i^*} \lambda_i(x) dx \\ &= \sum_{i \in \hat{B}} \int_{x_i}^{x_i^b} \lambda_i(x) dx + \left[ \sum_{i \in \hat{B}} \int_{x_i^b}^{x_i^*} \lambda_i(x) dx + \sum_{i \in B \setminus \hat{B}} \int_{x_i}^{x_i^*} \lambda_i(x) dx \right] \\ &\leq \sum_{i \in \hat{B}} \int_{x_i}^{x_i^b} \lambda_i(x) dx + \frac{1}{1-\alpha} \left[ \sum_{i \in \hat{B}} x_i^b \lambda_i(x_i^b) + \sum_{i \in B \setminus \hat{B}} x_i \lambda_i(x_i) \right] \\ &\leq \left( 1 + \frac{1}{1-\alpha} \right) \sum_{i \in \hat{B}} \lambda_i(x_i) x_i^b + \frac{1}{1-\alpha} \sum_{i \in B \setminus \hat{B}} x_i \lambda_i(x_i). \end{aligned}$$

The bound for the term in the square parentheses in lines two and three comes from the second half of Lemma B.3 for which we define a new demand vector  $\vec{x}'$  such that  $x'_i = \max(x_i^b, x_i)$ . The last line follows from simple rearrangement. The benchmark solution satisfies all of the conditions mentioned in Lemma B.6 and therefore, the first term above  $\sum_{i \in \hat{B}} \lambda_i(x_i) x_i^b$  is no larger than  $2 \sum_{t \in \hat{S}} p_t y_t + 2 \sum_{t \in \hat{S}} p_t^* y_t^*$ . Once again,  $\hat{S}$  is the set of saturated goods in our pricing solution.

In summary, we get

$$\sum_{i \in B} \int_{x_i}^{x_i^*} \lambda_i(x) dx \leq \left( 2 + \frac{3}{1-\alpha} \right) \sum_{i \in B} [\lambda_i(x_i) x_i + \lambda_i(x_i^*) x_i^*].$$

Coming back to our original expression, the difference in social welfare can be written as

$$\begin{aligned} SW^* - SW(\vec{p}, \vec{x}, \vec{y}) &\leq \left( 2 + \frac{3}{1-\alpha} \right) \sum_{i \in B} [\lambda_i(x_i) x_i + \lambda_i(x_i^*) x_i^*] + C(\vec{y}) \\ &\leq \left( 4 + \frac{6}{1-\alpha} \right) [\pi(\vec{p}, \vec{x}, \vec{y}) + \pi(\vec{p}^*, \vec{x}^*, \vec{y}^*)] + \pi(\vec{p}, \vec{x}, \vec{y}). \end{aligned}$$

The last step comes from Lemma C.3. □

### Penultimate Leg: Sequence of Solutions

Previously, we identified a suitable dummy price and proved that the solution obtained via the augmented equilibrium at this dummy price (to some extent) captures all of the welfare of the

benchmark solution and any lost welfare can be charged to the saturated goods. What about the profit of this solution? When all of the bundles desired by buyers are of equal cardinality, one can immediately show that the solution's profit also mimics that of the personalized payment scheme in the benchmark solution,  $\pi^b$ . When there is large disparity in the bundle sizes, however, the dummy price may result in buyers gravitating towards the smaller sized bundles. How do we fix this?

In this subsection, we consider augmented Walrasian equilibria at dummy prices that are scaled versions of the original dummy price  $\frac{\tilde{p}}{2^{\ell^{max}}}$ . Sequentially, we show that, if the profit at the previous solution (starting with the original dummy price) does not meet our requirement, then scaling the dummy price by a factor of two, does not really lead to a large loss in social welfare. Once again, our primary technique is a charging argument that assigns the lost welfare to the profit of the saturated goods. The profit due to the saturated goods is no larger than the social welfare of the solution, and thus, in either event, the social welfare cannot be small.

We begin with some notation. Consider the vector of prices defined as  $P(j) = 2^j \frac{\tilde{p}}{2^{\ell^{max}}}$  for  $j = 0$  up to  $j = \lceil \log(\Delta) \rceil + 1$ , and define  $(\vec{p}(j), \vec{x}(j), \vec{y}(j))$  to be the pricing solution for the original instance obtained via the augmented Walrasian equilibrium at dummy price  $P(j)$ . While below we assume that  $\Delta$  is a power of 2, this argument is easily generalized. Let us also define  $SW(j)$  and  $\pi(j)$  to be the social welfare and profit of the pricing solution  $(\vec{p}(j), \vec{x}(j), \vec{y}(j))$ . Observe that  $(\vec{p}(0), \vec{x}(0), \vec{y}(0))$  denotes the solution with the original dummy price for whose social welfare we obtained a bound in Claim 5.2.

Our next claim establishes a bound on the loss in social welfare as we increase the dummy price, in terms of the the previous and the current profit. For the rest of this proof, we define  $\delta := \log_2(\Delta)$ .

**Claim B.8** (Claim 5.3 from Section 5). *For every  $j \in [0, \delta]$ ,  $SW(j) - SW(j+1) \leq 3\pi(j) + 3\pi(j+1)$ .*

*Proof.* Consider  $SW(j) - SW(j+1)$ , this can be bound as follows,

$$SW(j) - SW(j+1) \leq \sum_{i \in B} \lambda_i(x_i(j+1))x_i(j) - (C(\vec{y}(j)) - C(\vec{y}(j+1))).$$

Let us divide  $B$  into two groups of buyers:  $B_1$ , which are the buyers satisfying  $\lambda_i(x_i(j+1)) \leq 2\lambda_i(x_i(j))$  and  $B_2$  which represent the buyers for whom  $\lambda_i$  has more than doubled. Our proof will proceed by charging the loss in welfare for the buyers in  $B_1$  to  $\pi(j)$ ; for the buyers in  $B_2$ , the loss in welfare will be charged to both  $\pi(j)$  as well as the income/revenue from the saturated goods in  $\pi(j+1)$ . In fact, we are only interested in a special class of saturated goods, for which  $c_t(y_t(j+1)) \geq c_t(y_t(j))$ . Formally, define  $SS(j+1)$  as the subset of saturated goods in  $(\vec{p}(j+1), \vec{x}(j+1), \vec{y}(j+1))$  having  $c_t(y_t(j+1)) \geq c_t(y_t(j))$ .

We first state a simple sub-claim, which we prove later.

**Lemma B.9.** *For every buyer  $i \in B_2$ , and any bundle  $S \in B_i$  that  $i$  consumes a non-zero amount of in  $(\vec{p}(j), \vec{x}(j), \vec{y}(j))$ , we have that  $\lambda_i(x_i(j+1)) - 2\lambda_i(x_i(j)) \leq \sum_{t \in SS(j+1) \cap S} p_t(j+1)$ .*

Now, we are ready to proceed with the proof. Let us divide  $\sum_{i \in B} \lambda_i(x_i(j+1))x_i(j)$  into two parts, namely (i):  $\sum_{i \in B_1} \lambda_i(x_i(j+1))x_i(j) + \sum_{i \in B_2} 2\lambda_i(x_i(j))x_i(j)$ , and (ii):  $\sum_{i \in B_2} [\lambda_i(x_i(j+1)) - 2\lambda_i(x_i(j))]x_i(j)$ . Consider the first part,

$$\sum_{i \in B_1} \lambda_i(x_i(j+1))x_i(j) + \sum_{i \in B_2} 2\lambda_i(x_i(j))x_i(j) \leq 2 \sum_{i \in B} \lambda_i(x_i(j))x_i(j).$$

We proceed with the second part carefully. Consider the solution  $(\vec{p}(j), \vec{x}(j), \vec{y}(j))$ , and define  $z_i(S)$  to be the amount of bundle  $S$  consumed by buyer  $i$  in this solution. Without loss of generality, assume that  $B_i$  consists only of bundles that  $i$  is consuming non-zero amounts of in  $(\vec{p}(j), \vec{x}(j), \vec{y}(j))$ . Applying our Sub-Lemma B.9, we get that,

$$\begin{aligned}
\sum_{i \in B_2} [\lambda_i(x_i(j+1)) - 2\lambda_i(x_i(j))] x_i(j) &\leq \sum_{i \in B_2} \sum_{S \in B_i} [\lambda_i(x_i(j+1)) - 2\lambda_i(x_i(j))] z_i(S). \\
&\leq \sum_{i \in B_2} \sum_{S \in B_i} \sum_{t \in SS(j+1) \cap S} p_t(j+1) z_i(S). \\
&\leq \sum_{t \in SS(j+1)} p_t(j+1) y_t(j) \\
&\leq \sum_{t \in SS(j+1)} p_t(j+1) y_t(j+1).
\end{aligned}$$

Finishing up the proof, we have that

$$\begin{aligned}
SW(j) - SW(j+1) &\leq 2 \sum_{i \in B} \lambda_i(x_i(j)) x_i(j) - C(\vec{y}(j)) + \sum_{t \in T} p_t(j+1) y_t(j+1) + C(\vec{y}(j+1)) \\
&= \pi(j) + \sum_{i \in B} \lambda_i(x_i(j)) x_i(j) + \sum_{t \in T} p_t(j+1) y_t(j+1) + C(\vec{y}(j+1)) \\
&\leq 3\pi(j) + 3\pi(j+1).
\end{aligned}$$

As usual, we replace the income with the profit as per Lemma C.3. Recall that  $\sum_{t \in T} p_t(j+1) y_t(j+1) = \sum_{i \in B} \lambda_i(x_i(j+1)) x_i(j+1)$ .

### Proof of Lemma B.9

We need to prove that for every buyer  $i \in B_2$ , and any bundle  $S \in B_i$  that  $i$  consumes a non-zero amount of in  $(\vec{p}(j), \vec{x}(j), \vec{y}(j))$ ,

$$\lambda_i(x_i(j+1)) - 2\lambda_i(x_i(j)) \leq \sum_{t \in SS(j+1) \cap S} p_t(j+1).$$

Consider the items in  $S \setminus SS(j+1)$ . Since  $\lambda_i(x_i(j+1)) = \sum_{t \in S} p_t(j+1)$  because  $i$  consumes a non-zero amount of bundle  $S$ , then in order to prove the lemma, it suffices to show that

$$\sum_{i \in S \setminus SS(j+1)} p_t(j+1) \leq 2 \sum_{t \in S \setminus SS(j+1)} p_t(j) \leq 2 \sum_{t \in S} p_t(j) = 2\lambda_i(x_i(j)).$$

We claim that for every  $t \in S \setminus SS(j+1)$ ,  $p_t(j+1) \leq 2p_t(j)$ . To see why this is true, assume by contradiction that  $p_t$  has more than doubled for some good  $t \in S \setminus SS(j+1)$ . Clearly, this good must be saturated in  $(\vec{p}(j+1), \vec{x}(j+1), \vec{y}(j+1))$  since the dummy price has only doubled, and  $c_t(y_t(j+1)) \geq 2c_t(y_t(j))$ . However, saturated goods which have  $c_t(y_t(j+1)) \geq c_t(y_t(j))$  belong to  $SS(j+1)$  by definition.  $\square$

Since we can bound the difference in welfare via profit, we can sequentially build solutions with good social welfare until we either reach a solution with good profit (in which case, we are done) or  $j = 1 + \delta$  is reached, and we have no more solutions to construct. In order to analyze the boundary condition, we now establish upper bounds on the social welfare at  $j = 1 + \delta$  in terms of the profit obtained at that solution.

**Lemma B.10.**

$$SW(\delta + 1) \leq \left( 2 \left( \frac{1}{1 - \alpha} \right)^{\frac{1}{\alpha}} - 1 \right) \pi(\delta + 1).$$

*Proof.* Consider the pricing solution  $(\vec{p}(\delta + 1), \vec{x}(\delta + 1), \vec{y}(\delta + 1))$ : this corresponds to an augmented Walrasian equilibrium with dummy price  $2^{\delta+1} \frac{\tilde{p}}{2^{\ell^{max}}} = \frac{\tilde{p}}{\ell^{min}}$ , where  $\ell^{min} \geq 1$  is the cardinality of the smallest bundle desired by any one buyer type.

Applying Lemma B.4, we know that for every  $t \in T$ ,  $p_t \geq \frac{\tilde{p}}{\ell^{min}}$ . Therefore, we have that for every buyer  $i \in B$ ,  $\lambda_i(x_i(\delta + 1)) \geq \frac{\tilde{p}}{\ell^{min}} |S|$ , for any  $S \in B_i$ . Since  $|S| \geq \ell^{min}$ , we conclude that  $\lambda_i(x_i(\delta + 1)) \geq \tilde{p}$ . Applying Claim B.3 on  $(\vec{p}(\delta + 1), \vec{x}(\delta + 1), \vec{y}(\delta + 1))$  completes the proof.  $\square$

Now, we are ready to combine all of the above claims and lemmas to make our final stand: a bound on  $SW^*$  in terms of the profits of various solutions, which would imply that the max-profit solution approximates the optimum welfare up to the desired bound. In the next subsection, we take this proof further and show that if the max-profit solution from the below claim does not result in a good welfare, then one can in fact identify another such solution with similar profit guarantees but *much better* welfare.

**Claim B.11** (Claim 5.4 from Section 5).

$$SW^* \leq \left( 8 + 2 \left( \frac{1}{1 - \alpha} \right)^{\frac{1}{\alpha}} + 4 \frac{1}{1 - \alpha} \right) \left[ \sum_{j=0}^{1+\delta} \pi(j) + \pi(\vec{p}^*, \vec{x}^*, \vec{y}^*) \right].$$

*Proof.* We know from Claim 5.2 that

$$SW^* - SW(0) \leq \left( 5 + \frac{6}{1 - \alpha} \right) [\pi(0) + \pi(\vec{p}^*, \vec{x}^*, \vec{y}^*)].$$

Our other crucial result, Claim 5.3 tells us that for any  $j$ ,

$$SW(j) - SW(j + 1) \leq 3\pi(j) + 3\pi(j + 1).$$

Writing the above inequality for all  $j \in [0, \delta]$  provides us with a nice telescoping summation, which upon addition with the first inequality above yields

$$SW^* \leq SW(\delta + 1) + \sum_{j=1}^{\delta} 6\pi(j) + 3\pi(\delta + 1) + \left( 8 + \frac{6}{1 - \alpha} \right) [\pi(0) + \pi(\vec{p}^*, \vec{x}^*, \vec{y}^*)].$$

Finally, by replacing  $SW(\delta + 1)$  with the bound obtained in Lemma B.10, we can wrap up the claim,



$$\begin{aligned}
SW^* &\leq (2 + 2(\frac{1}{1-\alpha})^{\frac{1}{\alpha}})\pi(\delta+1) + 6 \sum_{j=1}^{\delta} \pi(j) + \left(8 + \frac{6}{1-\alpha}\right) [\pi(0) + \pi(\vec{p}^*, \vec{x}^*, \vec{y}^*)] \\
&\leq \left(8 + 2(\frac{1}{1-\alpha})^{\frac{1}{\alpha}} + \frac{4}{1-\alpha}\right) \left[\sum_{j=0}^{\delta+1} \pi(j) + \pi(\vec{p}^*, \vec{x}^*, \vec{y}^*)\right]
\end{aligned}$$

For the second inequality, we used the fact that  $2\frac{1}{1-\alpha} \leq 2(\frac{1}{1-\alpha})^{\frac{1}{\alpha}}$ .  $\square$

## Final Pricing Algorithm

We have all of the required lemmas in our arsenal, and now will go about deriving our final algorithm. In the statement of Claim 5.4, there are  $2 + \log(\Delta)$  profit terms in the RHS; therefore, the solution giving maximum profit among  $\pi(0), \pi(1), \dots, \pi(\delta+1)$  and  $\pi(\vec{p}^*, \vec{x}^*, \vec{y}^*)$  must yield a  $(2 + \log(\Delta))\Theta(\frac{1}{1-\alpha})$  approximation to both profit and welfare. Our proof does not end here as our desired social welfare bound does not depend on  $\Delta$ , but by choosing the maximum profit solution, we may end up with a  $O(\log(\Delta))$  additional welfare degradation on our hands.

However, we show that by using some careful analysis, one can instead compute a solution whose approximation for social welfare is simply  $O(\frac{1}{1-\alpha})$  and thus, is independent of  $\Delta$ . We now define our main algorithm: for the purpose of continuity, let  $\pi(-1)$  represent  $\pi(\vec{p}^*, \vec{x}^*, \vec{y}^*)$ .

1. Let  $k$  be the smallest index in the range  $[-1, 1 + \delta]$  such that  $\frac{SW^*}{\pi(k)}$  is no larger than  $2(\log(\Delta) + 2)\left\{\left(8 + 2(\frac{1}{1-\alpha})^{\frac{1}{\alpha}} + 4\frac{1}{1-\alpha}\right)\right\}$ .
2. Return the solution  $(\vec{p}(k), \vec{x}(k), \vec{y}(k))$ .

From Claim 5.4, it is clear that there exists at least one index  $k$  providing the desired guarantee. In fact, if  $k = -1$ , then we are done because the solution returned is the one maximizing social welfare. The following claim provides a lower bound on the social welfare  $SW(k)$  when  $k \geq 0$ .

**Claim B.12.** *The algorithm provides a  $(\Theta(\frac{\log \Delta}{1-\alpha}), \Theta(\frac{1}{1-\alpha}))$ -bicriteria approximation to revenue and welfare.*

*Proof.* The revenue bound simply follows from the algorithm, so we focus on the social welfare here. Recall from Claim 5.2 that (i)  $SW^* \leq SW(0) + \left(5 + \frac{6}{1-\alpha}\right)(\pi(-1) + \pi(0))$ .

Moreover, applying Lemma 5.3 repeatedly for  $j = 0$  to  $k$  and performing the telescoping summation, we get that (ii) :  $SW(0) - SW(k) \leq 3\pi(0) + 6 \sum_{j=1}^{k-1} \pi(j) + 3\pi(k)$ .

We use the above inequalities to derive a lower bound on  $SW(k)$ . We prove this in two cases,

**Case I:**  $k = 0$

$$\begin{aligned}
SW^* &\leq SW(0) + \left(5 + \frac{6}{1-\alpha}\right)(\pi(-1) + \pi(0)) \\
\Rightarrow \frac{SW^*}{2} &\leq 6 \left(1 + \frac{1}{1-\alpha}\right) SW(0).
\end{aligned}$$

Clearly  $\pi(-1)(5 + \frac{6}{1-\alpha}) \leq \pi(-1)(8 + \frac{4}{1-\alpha} + 2(\frac{1}{1-\alpha})^{\frac{1}{\alpha}}) \leq \frac{SW^*}{2}$  as per the definition of the algorithm since  $k \neq -1$ . Moreover, we also used the trivial inequality that  $\pi(0) \leq SW(0)$ . In conclusion, we have that  $SW(k) = SW(0)$  is a  $\Theta(\frac{1}{1-\alpha})$ -approximation to the optimum welfare.

**Case II:**  $k > 0$ .

We add up the inequalities (i) and (ii) from above to get

$$\begin{aligned} SW^* &\leq SW(k) + 3\pi(k) + (8 + \frac{6}{1-\alpha}) \sum_{j=-1}^{k-1} \pi(j) \\ &\leq SW(k) + 3\pi(k) + \frac{SW^*}{2} \\ &\leq 4SW(k) + \frac{SW^*}{2} \\ \implies \frac{SW^*}{2} &\leq 4SW(k). \end{aligned}$$

The second inequality is a consequence of the fact that for every  $j \in [-1, k-1]$ ,  $SW^* \geq 2(\log(\Delta) + 2)\{(8 + 2(\frac{1}{1-\alpha})^{\frac{1}{\alpha}} + 4\frac{1}{1-\alpha})\pi(j)$ . We add this inequality for  $j = -1$  to  $k-1$ , which consists of at most  $k+1 \leq \log(\Delta) + 2$  terms.

So, the worst welfare is obtained for  $k = 0$ .

□

□

## C Appendix: Cost functions

The first two lemmas below are paraphrased versions of similar results in [3].

**Lemma C.1.** *Consider two pricing solutions  $(\vec{p}_1, \vec{x}_1, \vec{y}_1)$  and  $(\vec{p}_2, \vec{x}_2, \vec{y}_2)$  such that  $\vec{p}_1 \geq \vec{p}_2$ , and suppose that  $\vec{y}_1, \vec{y}_2$  are the minimum-cost allocations for the buyer demands  $\vec{x}_1$  and  $\vec{x}_2$  respectively over all feasible allocations. Then,*

1.  $\vec{x}_2 \geq \vec{x}_1$ , i.e., every buyer type's demand is higher under  $\vec{p}_2$  than under  $\vec{p}_1$ .
2. For all  $t \in S$ ,  $c_t(y_t^1) \leq c_t(y_t^2)$ .
3. For every buyer  $i$ ,  $r_i(z^{\vec{1}}) \leq r_i(z^{\vec{2}})$ .

**Lemma C.2.** *Consider two buyer demands  $\vec{x}^1$  and  $\vec{x}^2$  where  $\vec{x}^2 \geq \vec{x}^1$  and the corresponding min-cost allocations are  $\vec{z}^1$  and  $\vec{z}^2$ . Then, the following provides a bound for the difference in costs*

$$\sum_{t \in T} (C_t(z_t^2) - C_t(z_t^1)) \geq \sum_{i \in B} r_i(\vec{z}^1)(x_i^2 - x_i^1).$$

**Lemma C.3.** *Let  $(\vec{p}, \vec{x}, \vec{y})$  denote a valid pricing solution for a given instance such that for all  $t$ ,  $p_t \geq c_t(y_t)$ . Then*

1.  $\sum_{i \in B} \lambda_i(x_i)x_i \geq 2C(\vec{y})$ .
2.  $\pi(\vec{p}, \vec{x}, \vec{y}) \geq C(\vec{y})$ .

$$3. \sum_{i \in B} \lambda_i(x_i)x_i \leq 2\pi(\vec{p}, \vec{x}, \vec{y}).$$

*Proof.* The proof follows from the nature of doubly convex cost functions. Double convexity implies that for every  $t \in T$ ,  $C_t(y_t) \leq \frac{1}{2}c_t(y_t)y_t$ . Therefore, we have that

$$C(\vec{y}) \leq \frac{1}{2} \sum_{t \in T} c_t(y_t)y_t \leq \frac{1}{2} \sum_{t \in T} p_t y_t = \frac{1}{2} \sum_{i \in B} \lambda_i(x_i)x_i.$$

The final part of the above inequality comes from rearranging the allocation from the goods to the buyers, along with an application of the fact that for any buyer type  $i$  with non-zero consumption of bundle  $S$ ,  $\lambda_i(x_i) = \sum_{t \in S} p_t$ .

The second statement follows directly from the first since  $\pi(\vec{p}, \vec{x}, \vec{y}) = \sum_{i \in B} \lambda_i(x_i)x_i - C(\vec{y})$ . The third statement follows from the second.  $\square$

## D Appendix: Basic Properties of $\alpha$ -SR functions

**Lemma D.1.** *Let  $f(x)$  be any non-increasing, non-negative  $\alpha$ -Strongly regular function. Then for any constant  $c$ ,  $f(x) - c$  is also  $\alpha$ -strongly regular in the interval where it is non-negative.*

*Proof.* We need to prove that for any  $x_1 < x_2$ ,  $\frac{f(x_2)-c}{|f'(x_2)|} - \frac{f(x_1)-c}{|f'(x_1)|} \leq \alpha(x_2 - x_1)$ . We show this in two cases.

Case I:  $|f'(x_1)| \geq |f'(x_2)|$

When both of the derivatives are non-zero, we have that

$$\frac{f(x_2) - c}{|f'(x_2)|} - \frac{f(x_1) - c}{|f'(x_1)|} = \left\{ \frac{f(x_2)}{|f'(x_2)|} - \frac{f(x_1)}{|f'(x_1)|} \right\} - c \left( \frac{1}{|f'(x_2)|} - \frac{1}{|f'(x_1)|} \right).$$

Since the function  $f$  is  $\alpha$ -SR, the first quantity in the RHS is smaller than or equal to  $\alpha(x_2 - x_1)$  the second quantity is not positive as per the assumption for Case I. The lemma follows. Note that if any one of the two derivatives are zero, then the proof follows trivially.

Case II:  $|f'(x_1)| < |f'(x_2)|$

In this case,

$$\frac{f(x_2) - c}{|f'(x_2)|} \leq \frac{f(x_2) - c}{|f'(x_1)|} \leq \frac{f(x_1) - c}{|f'(x_1)|}.$$

The final inequality follows from the fact that the function  $f$  is monotone non-increasing. Therefore,  $\frac{f(x_2)-c}{|f'(x_2)|} - \frac{f(x_1)-c}{|f'(x_1)|} \leq 0 \leq \alpha(x_2 - x_1)$ .  $\square$

**Lemma D.2.** *Let  $f(x)$  be any non-increasing, non-negative  $\alpha$ -Strongly regular function. Then for any  $x_2 \geq x_1$ , the following inequality is true,*

$$\int_{x_1}^{x_2} f(x)dx \leq \frac{1}{1-\alpha} \frac{f(x_1)}{|f'(x_1)|} (f(x_1) - f(x_2)).$$

*Proof.* Consider any  $x \geq x_1$ . We now use the definition of  $\alpha$ -strong regularity to obtain a crude upper bound on  $f(x)$  in terms of  $f(x_1)$  and the derivative at  $x$ . i.e., we know that

$$\frac{f(x)}{|f'(x)|} \leq \frac{f(x_1)}{|f'(x_1)|} + \alpha(x - x_1).$$

For convenience, define  $h = \frac{f(x_1)}{|f'(x_1)|}$ . This gives us the following upper bound,  $f(x) \leq h|f'(x)| + \alpha|f'(x)|(x - x_1)$ . The quantity we wish to compute is  $\int_{x_1}^{x_2} f(x)dx$ . Using our upper bound from before, we have that

$$\int_{x=x_1}^{x_2} f(x)dx \leq h \int_{x=x_1}^{x_2} |f'(x)|dx + \alpha \int_{x=x_1}^{x_2} (x - x_1)|f'(x)|dx.$$

Recall that since  $f$  is non-increasing  $|f'(x)| = -f'(x)$ . Moreover, for the second integral, we can simply integrate by parts and get that  $\int_{x_1}^{x_2} (x - x_1)(-f'(x))dx = [(x - x_1)(-f(x))]_{x_1}^{x_2} - \int_{x=x_1}^{x_2} (-f(x))dx = -(x_2 - x_1)f(x_2) + \int_{x=x_1}^{x_2} f(x)$ . Combining the various pieces, we are now ready to complete the proof

$$\begin{aligned} \int_{x_1}^{x_2} f(x)dx &\leq h(f(x_1) - f(x_2)) - \alpha(x_2 - x_1)f(x_2) + \alpha \int_{x=x_1}^{x_2} f(x)dx \\ \implies (1 - \alpha) \int_{x_1}^{x_2} f(x)dx &\leq h(f(x_1) - f(x_2)) - \alpha(x_2 - x_1)f(x_2) \\ &\leq h(f(x_1) - f(x_2)). \end{aligned}$$

□

**Lemma D.3.** *Let  $f(x)$  be a non-increasing, non-negative  $\alpha$ -Strongly regular function and let  $x_1$  be any point satisfying  $\frac{f(x_1)}{|f'(x_1)|} \leq x_1$ . Then for every  $x \geq x_1$ , the following inequality holds,*

$$\frac{f(x)}{|f'(x)|} \leq x.$$

*Proof.* The proof follows almost directly from the definition of an  $\alpha$ -SR function. Since  $x \geq x_1$ , we have that

$$\frac{f(x)}{|f'(x)|} \leq \frac{f(x_1)}{|f'(x_1)|} + \alpha(x - x_1) \leq x_1 + \alpha(x - x_1) \leq (1 - \alpha)x_1 + \alpha x.$$

Since  $x_1 \leq x$ , we have that  $(1 - \alpha)x_1 + \alpha x \leq x$ . This completes the proof. □

**Lemma D.4.** *Let  $f(x)$  be a non-increasing, non-negative  $\alpha$ -Strongly regular function, and let  $x_1$  be any point satisfying  $\frac{f(x_1)}{|f'(x_1)|} > x_1$ . Then,*

$$f(0) < f(x_1) \left( \frac{1}{1 - \alpha} \right)^{1/\alpha}.$$

Notice that since  $f$  is non-increasing,  $f(0)$  is the maximum value of the function.

*Proof.* Consider the function  $g(x) = \log(f(x))$ . Differentiating this gives us,

$$g'(x) = \frac{f'(x)}{f(x)}.$$

Recall that since  $f(x)$  is non-increasing, its derivative cannot be positive. Integrating  $g'(x)$  from  $x_1$  to 0 gives us

$$\begin{aligned}\int_{x_1}^0 g'(x) &= \int_{x_1}^0 \frac{f'(x)}{f(x)} dx \\ g(0) - g(x_1) &= - \int_{x_1}^0 \frac{|f'(x)|}{f(x)} dx \\ \log \frac{f(0)}{f(x_1)} &= \int_0^{x_1} \frac{|f'(x)|}{f(x)} dx\end{aligned}$$

Now, since the function is  $\alpha$ -SR and  $x \leq x_1$ , we have that  $\frac{f(x)}{|f'(x)|} \geq \frac{f(x_1)}{|f'(x_1)|} - \alpha(x_1 - x)$ . The first quantity in RHS is at least  $x_1$  from the lemma statement, and therefore, we get  $\frac{f(x)}{|f'(x)|} > x_1(1 - \alpha) + \alpha x$ . Substituting this in the above integral, we get

$$\int_0^{x_1} \frac{|f'(x)|}{f(x)} dx < \int_0^{x_1} \frac{1}{x_1(1 - \alpha) + \alpha x} dx = \frac{1}{\alpha} \log\left(\frac{1}{1 - \alpha}\right).$$

Therefore, we get that  $\log \frac{f(0)}{f(x_1)} < \frac{1}{\alpha} \log\left(\frac{1}{1 - \alpha}\right)$ , giving us the upper bound of  $(\frac{1}{1 - \alpha})^{\frac{1}{\alpha}}$  for  $\frac{f(0)}{f(x_1)}$ .  $\square$

The following corollary that follows from the previous lemma is perhaps the most important ingredient required for our main theorem.

**Corollary D.5.** *Let  $f(x)$  be a non-increasing, non-negative  $\alpha$ -Strongly regular function and  $x_1$  be some point satisfying  $f(0) \geq (\frac{1}{1 - \alpha})^{\frac{1}{\alpha}} f(x_1)$ . Then,  $\frac{f(x_1)}{|f'(x_1)|} \leq x_1$ .*

*Proof.* Assume by contradiction that  $\frac{f(x_1)}{|f'(x_1)|} > x_1$ . Then as per Lemma D.4, it must be true that  $f(0) < f(x_1)(\frac{1}{1 - \alpha})^{\frac{1}{\alpha}}$ , which of course contradicts the corollary statement.  $\square$